







Emotion regulation contagion drives reduction in negative intergroup emotions

Received: 18 April 2024

Accepted: 17 January 2025

Published online: 06 February 2025


 Check for updates

Michael Pinus ^{1,2,3,7}, Yajun Cao^{1,2,7}, Eran Halperin ³, Alin Coman ⁴, James J. Gross ⁵ & Amit Goldenberg ^{1,2,6} 

When emotions occur in groups, they sometimes impact group behavior in undesired ways. Reducing group's emotions with emotion regulation interventions can be helpful, but may also be a challenge, because treating every person in the group is often infeasible. One solution is to treat a fraction of a group, and then hope the effect of the treatment will spread to other group members. To test the viability of this option, we designed an experiment to examine the impact of emotion regulation applied to different proportions of groups of six Israeli participants ($N = 2659$) who shared real-time responses to negative emotions-inducing political stimuli. Before interacting with each other, we treated different proportions of each group with an emotion regulation intervention called cognitive reappraisal, which involved teaching participants to reinterpret events to reduce negative emotions. The results showed that as the proportion of participants who received the treatment increased, there was a reduction in emotions within the non-treated participants. Furthermore, targeting above 40% of participants resulted in reliable group emotional change. Using semantic projection analysis, we validated the contagion of reappraisal language. These findings shed light on the conditions that enable collective emotion regulation.

Emotions are a central driver of human behavior, but they by no means always lead to desired outcomes. In many situations, achieving long-term individual and group goals requires the management of emotions via emotion regulation, defined as the activation of a goal to change the emotional trajectory¹. To facilitate successful emotion regulation, researchers have developed a variety of emotion regulation interventions. The most prominent of these is a reappraisal intervention, which involves changing how one thinks about a situation to influence one's emotional response^{1,2}. In reappraisal interventions, participants are taught to generate alternative interpretations of emotional situations. The advantage of reappraisal is that it is cheap, quick, and easy to explain. It is also effective, and seems to consistently help individuals regulate their emotions in a variety of contexts³⁻⁵.

One domain in which reappraisal interventions have been proven helpful is intergroup conflicts. Intergroup conflicts are characterized by negative intergroup emotions that contribute to hostility and violence⁶⁻⁸ and have a detrimental impact on the physical and mental well-being of millions around the globe⁹⁻¹¹. Reappraisal interventions have been successfully employed in intergroup conflicts. For example, in the context of the Israeli-Palestinian conflict, teaching Jewish Israelis how to use reappraisal led to a reduction in negative emotions and an increase in support for conciliatory policies toward Palestinians^{12,13}. These results were recently conceptually replicated in Colombia¹⁴. Porat and colleagues have further expanded the use of reappraisal by developing ReApp, an online application that gamifies reappraisal training for Israelis, showing that it reduced negative emotions towards the conflict¹⁵. But even in the examples above, in which

¹Harvard Business School, Boston, MA, USA. ²Harvard University, Digital, Data and Design Institute, Boston, MA, USA. ³Department of Psychology, Hebrew University of Jerusalem, Jerusalem, Israel. ⁴Department of Psychology, Princeton University, Princeton, NJ, USA. ⁵Department of Psychology, Stanford University, Stanford, CA, USA. ⁶Department of Psychology, Harvard University, Boston, MA, USA. ⁷These authors contributed equally: Michael Pinus, Yajun Cao.  e-mail: agoldenberg@hbs.edu

reappraisal interventions were found successful in improving inter-group relations, they targeted only small numbers of individuals. This is a limitation for employing emotion regulation interventions at scale because it is often impossible to conduct interventions on the whole group, both because of a lack of resources and the inability to reach every group member. Thus, it's critical to examine whether there is any "spillover" of effects from treated to non-treated participants, and if so, what the relationship might be between the proportion of treated participants and the impact on the whole group.

In the current project, we examined the possibility of leveraging the process of emotion regulation contagion to change group emotions (See Fig. 1). We define emotion regulation contagion as the spread of emotion regulation strategies from people who have received emotion regulation interventions to those who are not treated. We see emotion regulation contagion of cognitive reappraisal, which is tested in the current project, as driven by two processes (Fig. 1, right side). First, regulation contagion may occur as a result of social appraisals, which are defined as the adoption of other people's appraisal of a certain situation in a way that impacts one's emotional response^{16,17}. In the current case, non-treated participants would adopt the reappraisals of the treated participants to conflict situations. Second, we believe that regulation contagion is driven by social learning, such that non-treated participants would learn how to generate reappraisals to similar stimuli and apply this knowledge to new stimuli¹⁸⁻²⁰.

We wish to distinguish the process of emotion regulation contagion (in pink) from two similar processes (in blue, See Fig. 1). The first is emotion contagion, where the mere expression of emotions of the treated participants is impacting the emotions expressed by the non-treated participants^{21,22}. Although emotion contagion is likely playing a role in impacting non-treated participants' emotions, we examined whether non-treated participants are also adopting the use of reappraisal, in addition to adjusting their emotion expressions to match

those of other groups members. The second related process is compliance, where no real change in emotion experiences occurs, but merely a change in expressed emotions as a result of social pressure²³. Here again, although compliance may impact the results, we will show some evidence that non-treated participants are changing their emotions privately.

The literature on contagion in networks provides ample support for the idea that behavioral changes might spread from treated to non-treated participants^{24,25}. In the past few years, this idea has been expanded to the realm of psychological interventions, designed to elicit changes in various psychological processes^{26,27}. Previous research has primarily examined mainly how either the shape of the network^{28,29} or the location and type of seeded participants impact the spread³⁰⁻³². Here we test an unknown question about how many members of a group must be treated in order for the intervention to lead to group-wide emotion changes.

We used a text-based interaction design to precisely restrict the influence processes to written communication and to standardize the experimental setting across groups. Using the Israeli-Palestinian conflict as a case study, we examined how changing the proportion of participants completing a reappraisal intervention within small groups relates to negative emotion in both the treated and non-treated participants. This was done using a paradigm in which groups of six participants reacted to conflict-related images (by writing text and providing ratings) in real time, then saw each other's reactions. Before being exposed to the pictures and others' responses, we manipulated the proportion of participants within the group who completed a reappraisal intervention to range from zero and six participants. We then examined how our intervention impacted the negative emotions of the group as a whole as well as both the treated and non-treated participants' negative emotions.

Here, we show how the proportion of participants treated with a reappraisal intervention impacts the reduction of negative emotions

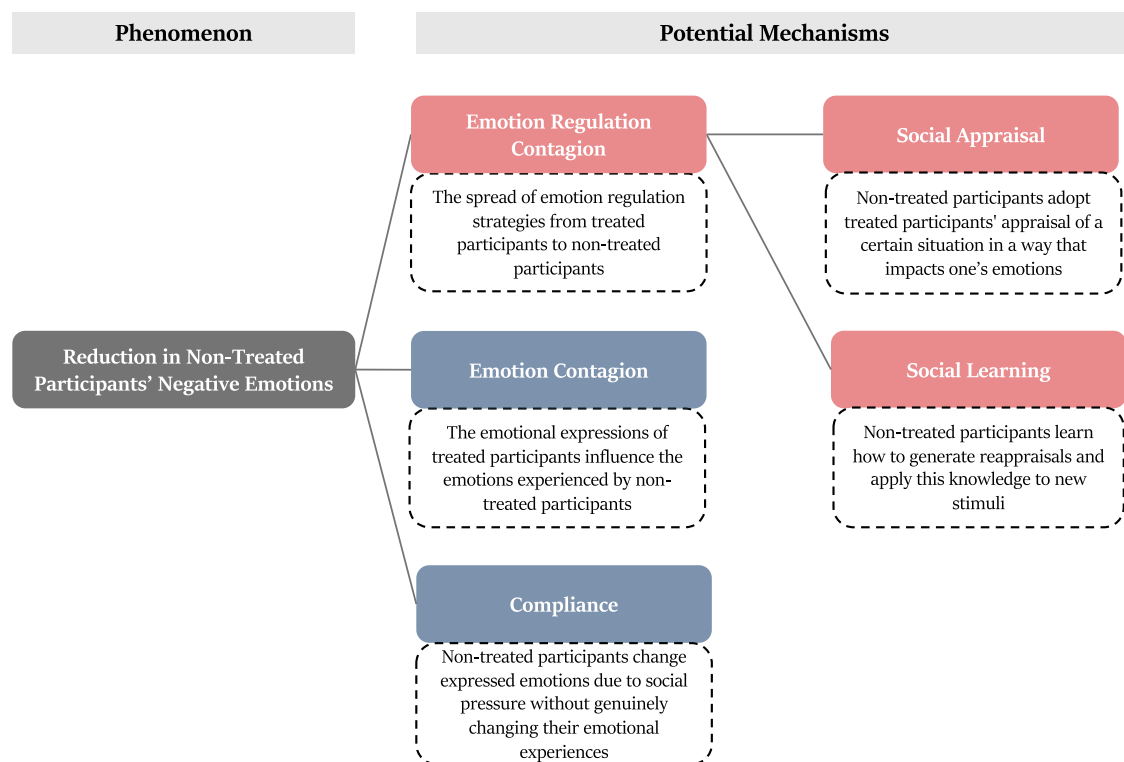


Fig. 1 | Potential mechanisms leading to a reduction in negative emotions among individuals not directly treated with emotion regulation interventions. The primary focus of this paper is emotion regulation contagion, which may

operate through social appraisal and social learning processes. Additionally, we aim to differentiate emotion regulation contagion from two related mechanisms: emotion contagion and compliance.

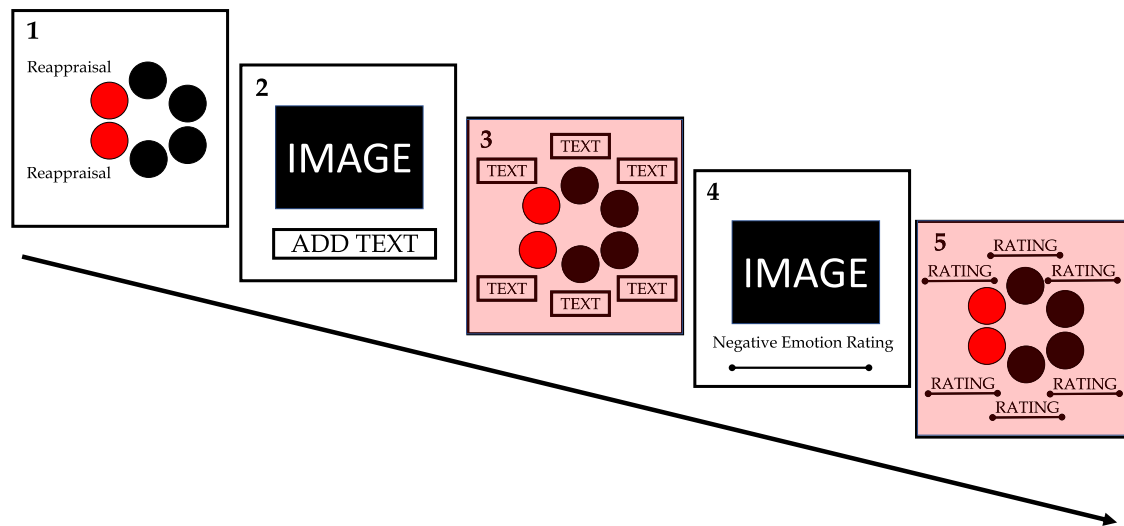


Fig. 2 | The structure of a trial in the emotion dynamics task (total of 20 trials). (1) Participants were assigned to groups of six and the proportion of participants who completed the reappraisal intervention was predetermined (in the example above, 2 out of 6 marked in red, see 1). (2) They then saw an image related to the conflict and were asked to provide their text to the picture. (3) Participants then saw

all the texts produced by everyone in the group. (4) They were then asked to rate their negative emotions to the picture from 1 to 10, (5) and then saw each other's ratings. There are two steps within the task (steps 3 and 5, which are shaded) in which people see each other's responses in real time.

within a group. We find that reappraisal reduces negative emotions not only for treated participants but also for non-treated participants. The relationship between the proportion of treated participants and reductions in non-treated participants' emotions is non-linear, and intervening on above 40% of participants resulted in reliable group emotional change. Semantic Projection Analysis³³ further reveals that non-treated participants adopt cognitive reappraisal strategies through social learning from treated individuals. An additional supplementary study eliminates the possibility that the reduction in emotion by the non-treated participants was driven solely by compliance.

Results

Israeli participants ($N=2659$) signed in to complete the study from their home computers. After consenting, participants were assigned to a group of six participants (Fig. 2). Before interacting with each other as part of our emotional dynamics task, a portion of the group (from zero to six participants) was assigned to a treated condition with a reappraisal intervention while the rest were assigned to a non-treated observing condition. Participants in the treated condition received instructions that were adapted from other reappraisal interventions³ and fitted to the Israeli context (see SI for full text). Participants were told that reappraisal is based on the insight that there are multiple interpretations for each situation, and that our emotional responses depend on these interpretations. They were then asked to practice reappraisal by reappraising a picture of an amputee meeting with a doctor and were then given examples of possible reappraisals to the situation. The observing condition was also used in previous tests of reappraisal interventions, and was found to lead to no significant changes in emotions compared to a passive empty control in which participants were not given any instructions about their emotions³. Participants in the non-treated condition were instructed to observe their emotions as they naturally unfolded. Similar to the intervention, they were also asked to practice observing their emotions by looking at the same pictures as in the reappraisal intervention. They were also given examples of possible emotional reactions to the situation.

The conditions in this study were randomly assigned in two steps. We first randomly assign the number of treated participants, varying between zero and six, to each group. Within each group, we then randomly assigned the individual condition (treated or non-treated) to

each participant. After being assigned to the group and to their condition (treated or non-treated), participants completed our emotional dynamics task in their group. During the task, participants saw pictures containing Palestinian resistance to the Israeli occupation and Palestinian violence against Israel. Pictures were mostly of terror attacks, or Palestinian demonstrations against Israel. These pictures were used in previous studies to elicit strong negative emotions among Jewish Israelis, mostly anger and sadness³⁴. After viewing each picture, participants were asked to produce a brief text that expressed their emotions (*"What comes up for you when you see the picture?"*, see further details in methods). Participants were then able to see the texts of all other participants in real time. They were then asked to rate their emotions in response to the picture on a 1-neutral to 10-very negative scale and again saw others' real-time ratings (*"Please rate the degree of negative emotions you feel in response to the picture"*, see full description in Methods). Participants therefore had two points during each trial when they could impact others' emotions, when observing others' text and when observing others' ratings (marked in red in Fig. 1). In total, there were 20 trials in this task.

It is important to note that due to natural dropout in online studies, the size of the group, and the actual proportion of people that went through the reappraisal manipulation, sometimes changed during the task. Therefore, the actual group size could be smaller than six. To counter these variations during the task, we used the actual proportion of reappraisers within each group and in each trial rather than the assigned proportion. We realize that it is often preferred to examine the assigned proportion for the analyses as it maintains random assignment. To make sure that results are consistent, we also conducted as-treated analyses which are based on the proportions of treated participants as originally assigned, and similar reduction in emotion by both the treated and non-treated participants as a function of proportion was found with this analysis (See SI). We also controlled for group size in all of the following models. After finishing the task, participants completed a survey that tested both manipulation checks such as intention to use reappraisal, and general sentiments towards Palestinians. These measures were designed to test whether there were changes in emotions that could be seen when participants know that their ratings will be shown to others, in order to reduce the possibility that compliance was driving participants' emotions. Using these more general ratings also allowed us to examine whether changes in

emotional ratings throughout the task, extended to more general sentiments toward Palestinians.

We conducted Kolmogorov–Smirnov tests on the normality assumption for regression analyses for the main hypotheses reported in the paper. When the assumption was violated, we conducted a robust estimation of mixed effects using the package `robustlmm`, finding similar results in all cases. All relevant statistical tests in the paper were two-tailed tests.

Before running the actual study, we conducted two pilot studies to validate key aspects for the analysis (see SI for full details). In the first pilot ($N=217$), we tested a Hebrew version of a reappraisal intervention³, examined among individuals and not in group contexts, and found the number of people it required to show a reduction in negative emotions to conflict-related stimuli. In the second pilot ($N=379$, see SI), we compared people's emotions in response to the stimuli either in groups of six – in a similar design to the one described above but with no people going through a reappraisal intervention – or when completing the task without being exposed to emotional responses of other group members. Results suggested that when participants were exposed to the stimuli in groups of six, but without having anyone assigned to the reappraisal intervention, they tended to express stronger emotions compared to when exposed to the stimuli separately, without seeing others' responses. Not only were participants' emotions stronger when seeing the stimuli in a group compared to separately, but their emotions also tended to intensify over trial numbers, suggesting a process of amplification over time. Finally, we examined emotion contagion by looking at changes in the variance of emotional ratings within the group over trial numbers. Results suggested that variance in emotions within the group decreased as the task progressed, providing evidence for contagion within the groups.

Was reappraisal effective for the treated individuals, even when some people in the group were not treated with reappraisal?

We first tested the effect of reappraisal by comparing negative emotion ratings as a function of whether participants were assigned to the treated or non-treated condition, across trials and regulators' proportions. We preregistered three ways to examine the effect of the intervention, and all of our tests were significant (See SI). Here we report the simplest way to examine whether the reappraisal treatment produced the hypothesized effect: a linear mixed model that predicted rating with the treated/non-treated condition, controlling for the actual proportion and the actual number of participants in the group in each trial. The model included random intercepts for the stimuli, the group, and the individual participants (nested within groups). The effect of the manipulation was statistically significant and negative ($t(2337.12) = -6.71$, $p < 0.001$, $\beta = -0.18$, 95% Confidence Intervals $[-0.23, -0.12]$), indicating that participants who were treated with the reappraisal intervention reported less negative emotions than the non-treated participants. We also tested the effect of reappraisal in a group setting by comparing participants' self-reported use of emotion regulation (see methods, $t(2335.51) = 44.09$, $p < 0.001$, $d = 1.88$, 95% Confidence Intervals $[1.78, 2.13]$), and the effort exerted on emotion regulation (see methods, $t(2551.57) = 39.26$, $p < 0.001$, $d = 1.51$, 95% Confidence Intervals $[1.45, 1.63]$). These effects were statistically significant and in the expected direction.

What is the relationship between the proportion of participants who went through the intervention and its effectiveness?

We hypothesized that higher proportions of treated participants in the group would lead to greater reduction in negative emotion, both within the treated and nontreated participants in each group. However, we had no specific prediction as to the shape of the reduction effect (i.e., linear or non-linear), so we compared alternative models to find the best approximation of the regulation agents' proportion dosage effect. We chose five different alternative models based on the

following rationale. The linear model was included to test for a simple, proportional relationship where emotions change consistently with the proportion of treated participants. The quadratic model was included to examine the possibility of an inverted U-shaped relationship, where emotions would initially increase with treatment in some form of reactance³⁵, but then reverse as the proportion continues to rise. The cubic model was considered to capture situations where contagion may be particularly potent either in low or high proportions. This could be driven by the way people represent the number of group members who are regulating, as previous research on social perception revealed a cubic model in the way people represent collective information³⁶. The logarithmic model was included to test the hypothesis that emotions might change rapidly at first as the proportion of treated participants increases but then taper off, indicating diminishing returns of the proportion of treated participants. Finally, the exponential model was chosen to explore the possibility that there is relatively little change in emotions when the proportion of treated participants is low, but the impact becomes significantly more pronounced as the proportion increases, suggesting a potential threshold effect.

Our first preregistered model was a linear mixed model in which we examined whether the proportion of reappraisal within the group (0%–100%) predicted negative emotions. The model included random intercepts for the stimuli, the group, and the individual participants (nested within groups). We fitted models representing different dosage levels in terms of the proportion of treated participants: linear, quadratic, cubic, logarithmic, and exponential. Results suggested that the strongest model was the exponential model ($t_{(885.01)} = -9.27$, $p < 0.001$, $\beta = -0.15$, 95% Confidence Intervals $[-0.18, -0.12]$) such that the reduction in emotions became increasingly larger with the increase in reappraisal proportion. As preregistered, in another model, we included an interaction term between proportion and condition and tested the simple effects in each condition. In that model, we found a statistically significant main effect of the exponential term of proportion ($t_{(2910.95)} = -3.19$, $p < 0.001$, $\beta = -0.10$, 95% Confidence Intervals $[-0.16, -0.04]$). The coefficient of the interaction term was not statistically significant ($t_{(3126.45)} = -3.19$, $p = 0.840$, $\beta = 0.001$, 95% Confidence Intervals $[-0.07, 0.08]$) (See SI for details).

An important limitation of the model described above is that it assumes the same dosage effect for both treated and non-treated participants. However, looking at the group as a whole may miss important information, as it seems possible that there might be different dynamics for each condition, and that the model would average over these differences. To account for this limitation, our preregistered analysis plan included performing the model comparison procedure described above for the treated and non-treated participants separately.

For the participants treated with reappraisal, all the models performed with similar AICs, but the model with the best fit was the cubic model, $AIC = 83,465.37$, which suggested that reduction in emotion was stronger within low and high proportions of reappraisal (see Fig. 3A). Within the cubic model, the effect of the proportion of the regulators was significant, ($t_{(621.17)} = -3.88$, $p < .001$, $\beta = -0.08$, 95% Confidence Intervals $[-0.12, -0.04]$). In a similar manner, in the non-treated subset of participants (i.e., those that were not presented with the reappraisal instructions), the model with the best fit was the quadratic model, $AIC = 11,082.4$ (see Fig. 3A). which suggested that there is relatively little change in the emotions of the non-treated participants until a certain proportion, at which point changes become much greater with every increase in the proportion of treated participants. Within the quadratic model, the effect of the proportion of the regulators was again significant ($t_{(1369.4)} = -3.46$, $p < 0.001$, $\beta = -0.06$, 95% Confidence Intervals $[-0.09, -0.02]$). It is worth mentioning that the AIC differences are minor, and the model comparison result should be taken with caution.

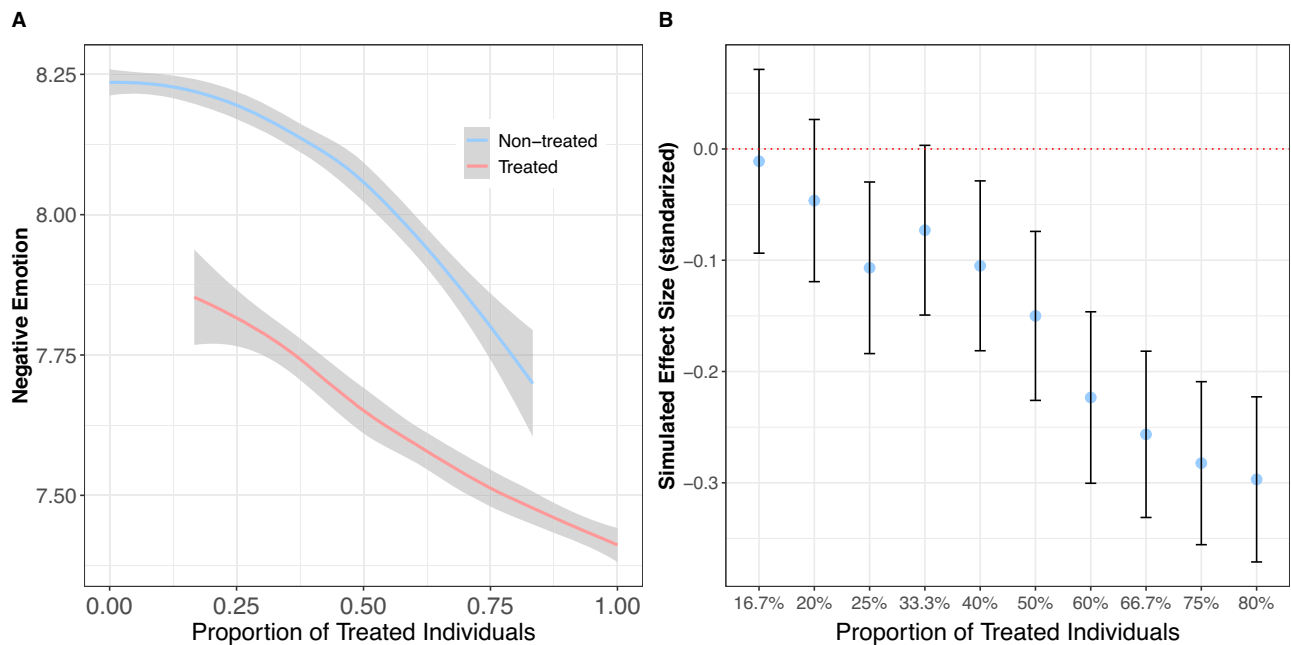


Fig. 3 | Results from the main study. Panel (A) captures the impact of proportion of regulators on negative emotions. The red line represents the estimated negative emotion ratings of the participants who were treated with reappraisal within each group of six ($N=1173$). The blue line represents the estimated negative emotion ratings of the non-treated participants within each group ($N=1486$). Grey areas represent standard errors. Results suggest that for the reappraisal condition, the best fitting model was a cubic model, although this model was very similar to others in terms of model fit. In the non-treated condition, the best fitting model was the quadratic model. Panel (B) captures the results of the simulation testing the

proportion of regulators needed to reach a significant change within the non-treated conditions. The blue dots and error bars represent the average standardized effects and their 95% critical intervals in 1000 iterations of simulation (simulated sample size $N=243$ per proportion condition; See details in Method). Using simulated data, we made sure that proportion bins included the same number of non-treated participants. We then compared the emotions of the groups with only non-treated participants to participants in the non-treated condition in each of the proportions. Results suggest a significant difference, with a reduction of 0.1 sd already at 25% proportion of regulators.

Overall, our results indicate that the dosage effect of the emotion regulation intervention is exponential at the entire group level. In addition, our results indicate that while the dosage effect might be different for the treated and non-treated participants, both subsamples demonstrated a degree of emotion regulation contagion, and showed non-linear reduction in conflict-related negative emotions.

Estimating when the impact becomes significant for the non-treated participants

One important question is what proportion of treated participants is required in a group for the non-treated participants in that group to be influenced by the reappraisal intervention. It is impossible to answer this question without data imputation with simulations, because in the raw data there are different numbers of reappraisal and non-treated participants in each proportion. The unbalanced number of reappraisal and non-treated participants in each proportion create unequal variances which make it difficult to compare effects. To equalize sample sizes for each proportion, we binned the data based on actual proportions of treated participants (see methods for detailed description). We then kept proportion bins that had more than 20 non-treated participants, and simulated the missing data to have 243 non-treated participants, which was the largest number of non-treated participants in each proportion. Imputation was done first by creating new groups for each proportion. Group size was determined based on the group size distribution within the task (see methods). We then populated these groups by simulating data based on the ratings of the participants in each proportion. The result of each simulation was groups of 243–245 non-treated participants that were equal in size for each proportion (numbers varied because of differences in group sizes). To test whether increasing the proportion of reappraisers led to a significant reduction in the ratings of the non-treated participants,

we compared the ratings of the non-treated participants in the baseline condition (0% reappraisal) to those of the non-treated in the different proportions of reappraisal. To make sure that our results were not driven by a specific simulation, we repeated the process 1000 times, each time comparing the non-treated participant only to all other conditions. In Fig. 3B we reported the standardized differences for these 1000 comparisons, with 95% confidence intervals. Results showed a reduction in non-treated participants' rating which became significant with a reduction of 0.1 SD in ratings already at 25% (Fig. 3B), and that treating 40% of the group with a reappraisal intervention results in significant and reliable emotion contagion at the group level. With the largest proportion (80%) in the sample, non-treated participants' negative emotions were reduced by nearly 0.3 SD. While this very much depends on the size of the bins, it provides a sense of comparison for future studies.

Impact on general sentiments towards Palestinians

After completing the task, participants were asked to complete a survey of more general sentiments towards Palestinians. We use the term sentiments, inspired by Frijda's conceptualization³⁷, because they do not reflect emotional responses to a specific situation but rather more general feelings towards Palestinians. Unlike the ratings during the task, participants knew that their ratings to these sentiment questions were not going to be shown to others. This served as a good opportunity to examine changes in emotion without the peer pressure of having others view their ratings. We reasoned that finding differences in the expression of sentiments towards Palestinians would be another indication that genuine contagion occurred. Note that we decided not to have a measure of general sentiments towards Palestinians before the task because we did not want such a measure to impact the quality of reappraisal training.

Participants were asked to rate nine negative sentiments towards Palestinians (e.g., “Generally speaking, when you think about the Palestinians in the Palestinian territories, to what extent do you feel fear towards them?”). Each item was rated on a 1–6 scale (1 - not at all, 6 - very much). We created a negative emotional attitudes scale using the relevant emotional items (fear, anger, hatred, disgust, $\alpha=0.81$). We reported results on positive emotions and guilt in supplementary information.

To examine the effect of the treatment on the sentiments towards Palestinians following the task we averaged both the proportion of treated participants and the group size across all trials. We then conducted a mixed model interaction between the proportion of participants treated with reappraisal in the same group and the condition assigned to the specific participant predicting the different sentiments. For the proportion of treated participants, we used an exponential model, as this was suggested to be the best fitting model for the interaction, but results were similar with a cubic or linear model (see SI). Our model also included a random intercept of group as well as condition nested within group, as participants were nested within different groups. Looking first at negative emotions and exploring the main effects, results suggested that increasing the proportion of participants treated with reappraisal led to lower negative sentiments for both treated and non-treated participants ($t_{(1772.34)} = -2.50$, $p = 0.01$, $\beta = -0.13$, 95% Confidence Intervals [-0.24, -0.04]). There was not a significant main effect between the treated and non-treated participants when ignoring the proportion of treated participants ($t_{(1767.21)} = -0.80$, $p = 0.42$, $\beta = -0.02$, 95% Confidence Intervals [-0.08, 0.06]). However, we did find an interaction between the proportion of treated participants and condition, such that the relationship between proportion of treated participants and negative sentiment was stronger for the non-treated participants ($t_{(1227.51)} = 1.97$, $p = 0.04$, $\beta = 0.10$, 95% Confidence Intervals [0.01, 0.21]).

Overall, these results suggest that the effect of proportion of treated participants on emotional ratings was also extended to negative general sentiments towards Palestinians. These results are encouraging because participants provided these ratings knowing that no other participants would see them. It is therefore another support that the manipulation led to real changes in emotion. It’s worth mentioning that we also measured general attitudes as well as dehumanization towards Palestinians. Results pointed to significant reduction in dehumanization (which are closely related to negative sentiments of anger, hate, and contempt) and marginally significant reduction in negative attitudes towards Palestinians, but as expected these results were weaker than the emotional results (see SI).

Providing evidence for the spread of reappraisal in semantic content

Our results provide evidence for reduction in emotion as a result of the increased proportion of participants treated with reappraisal, but we have not yet provided evidence that reduction in ratings within the non-treated participants is driven by changes in their interpretation of the situations. We sought to show the changes in participants’ appraisals by examining changes in the text that they produced as a function of the proportion of participants treated with reappraisal. To do this, we utilized a method called Semantic Projection Analysis³³, which is based on the idea that semantic meanings can be estimated by subtracting one linguistic representation from another. For example, to create a semantic representation of the term Queen, one can take a semantic representation of the term King and subtract the difference between the semantic representation of the term man and woman. Using the same idea, to generate a linguistic representation of a “pure reappraisal” content we can take the content produced by participants who were assigned to exclusively reappraisal groups (i.e., all 6 participants in the group were

treated with the reappraisal intervention) and subtract the content produced by participants who were assigned to exclusively non-treated groups (i.e., none of the 6 participants were taught to reappraise). The result is a semantic representation of “pure reappraisal”. We can then compare this pure reappraisal representation to the texts that participants produced throughout the task. At this point it is important to acknowledge that some non-treated participants may spontaneously reappraise, and some participants treated with reappraisal may not reappraise (despite their instructions to do so). This means that any findings from this analysis must rise above this noise.

To conduct our Semantic Projection Analysis, we processed the text to derive a 768-dimension embedding vector for each text response, using AlephBERT, a large pre-trained language model for modern Hebrew³⁸. Next, we created aggregated baseline vectors for the treated and non-treated response, by selecting only the responses that were provided by participants who were in groups that were pre-allocated to either 0% or 100% regulators. We then subtracted the non-treated baseline vector from the regulation baseline vector, to derive the “pure reappraisal” vector. We then computed the cosine distance of each individual’s text responses in our dataset to the “pure reappraisal” vector, to estimate the usage of reappraisal language in it. Lastly, we fitted a mixed-linear model to predict the usage of reappraisal language as a function of the proportion of reappraisers in each group. More specifically, we conducted a three-way interaction between condition (treated or non-treated), the proportion of reappraisers in the group (0%–100), and the trial number (see Fig. 4B). Similar to previous models, we used random intercepts for stimuli, group, and individual participant (nested within groups).

Focusing on the main effects, results showed that regardless of the proportion of treated participants, those who were assigned to the treated condition (and thus exposed to the reappraisal intervention) within each group of 6 were more similar to the “pure reappraisal” content than the non-treated condition, ($t_{(1220.79)} = 13.76$, $p < 0.001$, $\beta = 0.35$, 95% Confidence Intervals [0.30, 0.40]). This was expected given the assigned conditions and served as a sanity check. Results also indicated a main effect of proportion: increase in the proportion of participants in the group who were taught to reappraise led to an increase in similarity to the “pure reappraisal” semantic representation (Fig. 4A; $t_{(2060.14)} = 3.46$, $p < 0.001$, $\beta = 0.08$, 95% Confidence Intervals [0.03, 0.12]). Finally, we also found a significant effect of time, such that an increase in trial number led to an increase in similarity to the “pure reappraisal” semantic representation, ($t_{(28097.55)} = 10.57$, $p < 0.001$, $\beta = 0.07$, 95% Confidence Intervals [0.06, 0.08]).

Having established these three main effects, we then examined interactions. The only significant interaction was that between condition (treated or non-treated) and trial number ($t_{(28098.55)} = -7.65$, $p < 0.001$, $\beta = -0.08$, 95% Confidence Intervals [-0.10, -0.06]), suggesting that the association between trial number and similarity to “pure reappraisal” was stronger for the non-treated condition than the treated condition (Fig. 4A). This finding emphasizes that the non-treated participants were much more influenced by the reappraisal language as a function of the proportion of reappraisers than the treated participants.

To further investigate the relationship between proportion of treated participants and semantic similarity to “pure reappraisal,” we examined the simple effect of each condition separately. Results suggested that increasing the proportion of reappraisal led to a significant increase in similarity to pure reappraisal for the non-treated condition ($t_{(951.21)} = 3.72$, $p < 0.001$, $\beta = 0.08$, 95% Confidence Intervals [0.03, 0.10]) and a marginally significant increase in the treated condition ($t_{(824.57)} = 1.98$, $p = 0.05$, $\beta = 0.05$, 95% Confidence Intervals [0.03, 0.10]). These results indicated that the spread of reappraisal language could be a function of increases in the proportion of reappraisers within each group.

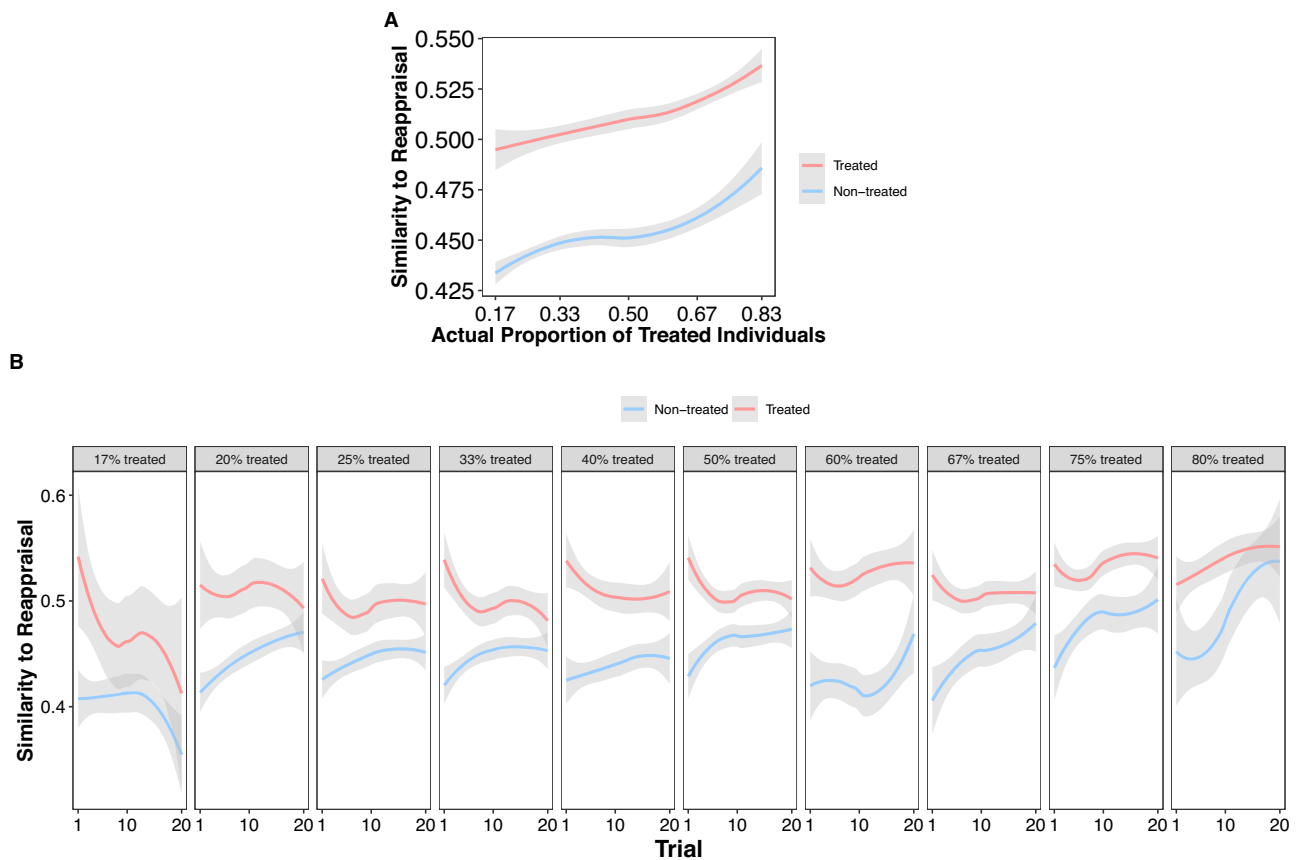


Fig. 4 | Semantic analysis of the text responses. **A** Similarity to reappraisal, evaluated using semantic projection analysis, for both the treated ($N = 1173$) and non-treated conditions ($N = 1486$). The x axis represents the proportion of participants treated with reappraisal in each group of six. The y axis represents the semantic similarity to the “pure reappraisal” semantic representations. Grey areas represent standard errors. Results suggest that increase in the proportion of participants who were treated within each group of six led to a marginally significant increase in similarity in the treated condition (red line) and to a significant increase

in similarity to reappraisal in the non-treated condition. **B** Similarity to reappraisal over time as binned by the proportion of reappraisers in each group, where the x axis represents the trial number (1–20), and the y axis represents the semantic similarity to the “pure reappraisal” semantic representations. Grey areas represent standard errors. Results suggest that when the number of treated individuals is low (17%) participants become more distant from reappraisal over time. However, as the number of treated individuals increases within the group, we see an increase in similarity to reappraisal language over time.

Separating the effect of groupmates’ texts from the negative emotion ratings

The fact that participants saw others’ texts and ratings presents two challenges to an interpretation of findings as supporting the idea of regulation contagion. First, it is possible that participants conformed to the ratings and not the texts, suggesting merely a simple contagion effect. Second, it is possible that participants didn’t even agree with the ratings of other group members, but merely changed their own ratings for the sake of compliance. We, therefore, aimed to design a study in which ratings and texts were separated to show that emotion regulation contagion is not solely driven by the influence of ratings and goes beyond the compliance effect. To do so, we conducted a supplementary study with a similar experimental process where participants responded to groupmates’ texts without seeing the ratings in their groups (See SI for details). In this study, we compared the negative emotions of participants who responded to all non-treated groupmates’ texts with those who saw texts from both treated and non-treated groupmates. This design allowed participants to give their ratings privately to mitigate the compliance effect and could also tease out the influence of groupmates’ texts from ratings. Results showed that those who viewed treated groupmates’ texts reported significantly less negative emotions than those who only saw non-treated groupmates’ texts, $t = -5.77$, $p < 0.001$, $\beta = -0.28$, 95% CI $[-0.37, -0.18]$. The finding from this supplementary study suggests that emotion regulation contagion existed independent

of the influence of others’ ratings, including the compliance effect driven by the public ratings.

Discussion

In the current study, we examined how the proportion of participants treated with a reappraisal intervention impacted the reduction in negative emotions within the group. More specifically, using inter-group conflict as the context for the investigation, we designed a paradigm in which groups of six participants responded emotionally to conflict-related stimuli. We then tested how increasing the proportion of participants treated with reappraisal impacted the emotions of the non-treated participants.

We found that reappraisal reduced participants’ emotions, even if the non-treated participants themselves were not instructed to use reappraisal. We also found that when the proportion of treated participants was above 40% (using the specific simulation we applied), there was a reliable difference in emotion ratings (compared to ratings of non-treated participants in groups with only non-treated participants). Analyzing participants’ text using Semantic Projection Analysis provided evidence for change in language produced by the non-treated participants as a function of the proportion of treated participants, providing support for linguistic contagion between the treated and non-treated participants. The linguistic analysis results provide evidence for the social learning mechanism, such that the non-treated participants seemed to learn the cognitive reappraisal strategy

from other treated individuals and use it in their newly generated texts. However, our limited findings from the groups in which they were only a small proportion of treated participants also point to the fact that the participants may not always be able to lead the group to adopt using reappraisal. As panel B of Fig. 4 shows, when only 1 out of 6 members in the group was a regulator, the similarity to reappraisal language in the texts of both treated and non-treated participants appeared to decrease over time. These findings are aligned with complex contagion theory such that the success of complex contagion depends upon interaction with multiple sources of affirmation^{29,39}.

Emotion regulation interventions designed to improve intergroup relations aim to impact the collective as a whole. Given limited resources and the fact that it's rarely possible to access all members of a certain group, processes of emotion regulation contagion are crucial for ensuring that interventions targeting individuals can influence the entire group. The current project turns the focus towards collective-level outcomes⁴⁰. This focus on the collective leads to a completely different set of questions. For example: How many people need to be treated with an intervention in order to achieve an overall outcome? Who should be targeted for such interventions for maximum efficiency? We believe that considering these questions could improve the utility of emotion regulation intervention but also other, more general psychological interventions. A variety of interventions, such as growth mindset⁴¹, social belonging⁴², and health behaviors⁴³, could benefit from thinking about their impact at the collective rather than individual level. We hope that this project is a significant step in a broader examination of the spread of psychological interventions within groups and collectives.

Limitations and future directions

This project has limitations related to the interpretation of the findings and the translation of these findings to applied contexts. The first limitation is related to the differentiation of various mechanisms behind emotion regulation contagion. It is possible that what led to the reduction in ratings was both emotion contagion, where mere exposure to emotion led to non-treated reduction in emotions, and emotion regulation contagion, where the actual use of reappraisal spread to the non-treated participants. Our text analysis provides a partial response to this issue. Non-treated participants who were exposed to responses of treated participants were more likely to use texts that were similar to reappraisal in their own responses. This is especially striking given that participants' text responses were provided to novel pictures, before seeing others' responses. Future research should try to isolate emotion contagion and emotion regulation contagion in more controlled ways.

A second limitation is the concern that changes in ratings throughout the task may have been driven by compliance with other group members, rather than real changes in emotion. We directly addressed this concern with a supplementary study where participants gave ratings privately showing changes in their ratings despite not being seen by other group members. Note that we do not think this study rules out the possibility that compliance may be playing some role in shaping participants' ratings. Instead, we believe that it provides strong support for the idea that compliance can't be the only factor driving performance. Additionally, participants' general sentiments towards Palestinians after the task, when participants knew that their ratings would not be observed by other participants, were also reduced with the increased proportion of treated participants, providing another evidence that compliance was not the only factor driving the results. Nevertheless, future studies should use other non-self-report methods to examine changes in emotion as a result of regulation contagion (see for example⁴⁴).

A third limitation relates to the fact that the size and nature of the groups – as well as the way in which they communicated – are different from typical groups, making it hard to generalize to natural interactions. Groups often vary in terms of size, hierarchy, and network typology, and

all of these aspects might influence how emotion regulation spreads. For example, it is likely that targeting central nodes in a network, or people with high power, may change the relationship between the number of treated participants and their impact on those who were not treated. We chose to simplify these aspects to reduce the noise and examine these contagion effects in a simple and controllable design. However, future studies should not only vary these group features, but should also try to examine the spread of interventions in natural field experiments. On a related note, social interactions in the current study were solely based on text communication. We believe this communication form can make our findings particularly relevant for digital contexts such as social media, but we also recognize that the absence of non-verbal information such as facial expressions and body postures may shape the contagion process of emotions and emotion regulation. Besides, in more natural conversations, reappraisal training might not only change how group members react to stimuli but also how people respond to each other's comments and the quality of discourse. Future studies should examine regulation contagion in more naturalistic environments such as face-to-face interactions in work teams⁴⁴.

A fourth limitation of the paper exists in its context of intergroup conflicts where people's group identities were particularly salient. This might influence people's willingness to take in social information and to express strong emotions in the group^{45,46}. In other contexts, where people's group identity was less salient, we might expect a different threshold of the proportion of treated group members for emotion regulation to work.

Finally, a fifth limitation of the current project is that it focused on reappraisal but did not compare its impact to another emotion regulation strategy. Furthermore, participants in the study were instructed to use the strategy on their own emotions, whereas in many situations people have an explicit goal to influence the emotions of others. Future studies should examine how different emotion regulation strategies and goals influence the relationship between proportion of treated participants and impact on the whole group.

Despite these limitations, we believe that the current project represents an exciting step for research on emotion regulation interventions and for psychological interventions in general.

Methods

The pilot and main study were approved by both the Harvard University Institutional Review Board (IRB) and The Hebrew University of Jerusalem IRB. The supplementary study was approved by the Harvard University IRB. Informed consent was obtained from participants in all studies. All analyses were conducted in R (Version 4.4.1). Preregistration documents can be found at <https://osf.io/d7u4h>.

Participants

Our power analysis for the main study was based on a pilot study in which we examined the number of participants required to achieve a significant difference between the treated and non-treated condition. We conducted a pilot study in which 217 participants were assigned to either a reappraisal intervention or a non-treated condition (see SI for full description). Results suggested that the effect size of the reappraisal intervention was $d = 0.20$, which suggested that 200 participants would be needed to detect a difference between the two conditions in terms of negative emotions. Because we did not have a good estimate of the expected size for the effect of the proportion of treated participants, and because we realized that the study may require more sensitivity, we decided to double the estimate and set the planned sample to 400 participants per condition, with total of 2800 participants.

Our aim was to recruit 2800 participants, 400 in each proportion group. Participants were Jewish Israelis who were recruited through iPanel, an Israeli survey company in exchange for 25 NIS (- \$7). Participants could start the task only after being grouped with five other participants. If after 5 minutes of waiting we were not able to group

participants, they were transferred to complete a different task. Therefore, our initial sample was much larger than the one we were able to group. Our initial recruited sample was 6161 participants. Of those participants, 5040 were able to be grouped in the initial round and started the task, but may have dropped in many different stages. The list of 5040 included participants who started the task twice. We allowed participants who did not start the actual rating phase (1604) to run through the task again. As preregistered, our dropout criteria were (numbers may be overlapping): (1) Participants who failed the reading check in which they were asked to choose the option that best describes the strategy that they were assigned to before starting the actual task (146 participants). (2) Participants who did not provide text responses, provided nonsensical text responses, or provided the exact same text response (10 participants). (3) groups who reported that they did not see the images or who had technical issues (6 groups). Participants were also automatically dropped from the study if they did not complete 15 of the 20 trials. After all of these omissions we had 2830 who completed the whole task. We removed all participants who were the only person remaining in the group, leaving us with 2659 participants which is our final sample. (Gender: 1158 males (43.6%), 1495 females (56.2%), 6 other or refused to say (0.2%); Age: $M = 42.07$, $SD = 14.618$). It is worth mentioning that the average attrition rate of non-treated participants (38.0%) is lower than that of the treated participants (50.2%). The difference in attrition rate might inflate the results because the treated participants may be generally better at reappraisal. Nonetheless, it is important to note that we are using the actual proportion of treated participants in our analysis at any time. This means whenever the dropout happened, it also influenced the proportion of the treated participants. We also conducted as-treated analyses with the original proportion values to ensure the robustness of our findings (See SI).

Task

When logging in to the task, participants were told that they were going to do a study in real time with 5 other participants. Participants were asked to choose a name and were told that other participants would see that name during the task. Following this stage, participants were forwarded to a waiting room where there were assigned to a group of six participants. Once six people logged in, the group was assigned a condition (0-6 people reappraising) and each person in the group was assigned to either the treated or non-treated condition. Both conditions were based on a recent reappraisal intervention that was validated in a large global sample³, and in a pilot study as a preparation for the current project (see SI). In the non-treated condition, participants were told that they would be asked to implement a strategy called *observing*, which involves paying attention to emotions as they unfold. We chose to use this active control condition – in which participants were asked to engage with their emotions – because it has been used in previous studies exploring reappraisal. It is, however, worth mentioning that a recent study that compared observing to a more passive non-treated condition in which participants were merely instructed to respond did not find consistent differences between the two conditions³. This can be a limitation particularly in the context of a conflict where having people reflect on their emotions may lead to stronger emotions. Participants in the reappraisal condition received instructions that were similar to Wang et al., but were slightly modified to the Israeli context (see SI for full text). Participants were told that reappraisal is based on the insight that there are multiple interpretations for each situation, and that our emotional responses depend on these interpretations.

Participants in both conditions then observed a practice image an amputee meeting with a doctor who is holding a prosthetic limb. They were then asked to respond to the picture and saw example responses based on the condition. Participants then completed an open ended question in which they were asked to describe the instructions of the

task as well as answered a multiple choice question in which they were asked to select the description of their condition. We removed participants who failed to properly answer both of these questions. Notice that the example given to participants was very different from the pictures in the actual study. This was done to avoid linguistic copying as much as possible from the practice stage to the actual task.

Finally, before the start of the task, participants were shown two pictures. For one picture (a picture of a truck) participants were told that the average negative emotion that the picture elicited was one, for the second (a picture of a child corpse) participants were told that the average rating was 10. The reason we added these descriptions is because in initial piloting we found that participants' responses to the pictures were almost at ceiling and we wanted to reduce the average emotional rating.

The actual task was 20 trials long. In each trial, participants were synchronously presented with a picture related to the Israeli-Palestinian conflict. Pictures were chosen from a sample of pictures used in a previous study because they elicited strong negative emotions, especially anger, among Jewish Israelis⁴⁷. Each picture contained a location, to clarify where they were taken. While observing the picture, participants were told: "Try to use the method you learned, observing/reappraisal, and to express your emotional response to this picture. The response should be short (one sentence). What comes up for you when you see the picture?" Participants were asked to enter a text to the picture and had up to 35 s to do so. Following this stage, participants were forwarded to a window in which they saw the name of the person responding (names were selected by participants in the beginning of the task) and their text in response to the picture. Participants were able to observe each others' responses for 15 seconds. Following this stage, participants were asked to rate their negative emotions in response to the picture on a scale of 1-no negative emotion to 10- very strong negative emotion. Participants had 25 s to rate their emotional responses and following this stage they again saw each user's name and their rating to the picture for another 15 s. Participants completed 20 trials of the task. After 10 trials, participants received a reminder of their instructions for the task (either reappraisal or control). Following the task, participants completed a few questions about demographics as well as a few exploratory surveys (see SI) that were mainly designed to examine potential mediators and were not preregistered in the analysis.

The experiment was conducted in multiple runs during April 2022. In every run, participants were randomly assigned to a group of 6 participants and to one of the 7 conditions, corresponding to the proportion of participants trained with reappraisal.

Measures

In addition to the measures taken in the task which are described above, participants also completed a survey following the task. The survey included three parts. The first part measures aspects related to emotion regulation. The second measured participants' sentiments towards Palestinians. The third measured attitudes towards Palestinians (described and reported in SI).

Emotion regulation. Participants were first asked about their emotion regulation attempts using a three item scale that was adopted from questions used in a previous reappraisal intervention⁴⁸ ($\alpha = 0.87$): "to what extent (if any) did you try to control your emotions while watching the pictures", "To what extent did you try to reduce negative emotions that came up while watching the pictures", "while watching the pictures, how much effort did you put to regulate your emotions?". Responses were rated on a scale of 1-not at all, to 6-very much so. Participants also rated the degree to which they used reappraisal in the task using a four item scale that was adapted from the same intervention reported above⁴⁸ ($\alpha = 0.84$): "While watching the pictures I tried to change their meaning.", "While watching the pictures I tried to give them a more positive meaning", "While watching the pictures I tried to understand

why people do what they do.”, “While watching the pictures I tried to give them a new meaning.” Responses were rated on a scale of 1-not at all, to 6-very much so. In addition to these two measures participants also completed the Emotion Regulation Questionnaire (ERQ)⁴⁹

General negative sentiments towards Palestinians. In addition to the measures in the task, we examined participants’ general sentiments towards Palestinians. We measured four negative emotions (fear, anger, hatred, disgust, $\alpha = 0.81$), and additional 4 positive emotions and guilt (reported fully in SI). For each emotion, participants were asked: “Generally speaking, when you think about the Palestinians in the Palestinian territories, to what extent do you feel [emotion] towards them?”. Each emotion was rated on a 1-6 scale (1-not at all, 6 – very much). We created a negative emotions scale (fear, anger, hatred, disgust) by averaging all of the negative emotion items.

Simulation

We conducted a simulation analysis to estimate when the impact of reappraisal becomes significant for the non-treated participants. The goal of the simulation is to mitigate the statistical bias caused by the unequal number of non-treated participants under different conditions. In our experiment sample, conditions with lower non-treated proportions had fewer observations of non-treated participants’ ratings. This would lead to higher variances in statistical estimates when we compared the reappraisal effect for non-treated participants across different proportion conditions. To make the reappraisal effects in all conditions comparable, it is necessary to make sure that these conditions have an equal number of observations for non-treated participants.

Data generation. We populated the non-treated participant sample by creating new groups and simulated non-treated participants’ ratings in each group. We first needed to decide on the group size of each simulated group as the group sizes in the actual experiment varied across groups due to dropouts. We sampled the group size data from a normal distribution with the mean and standard deviation of the group sizes of all existing groups in the original sample ($M = 3.72$, $SD = 1.10$). We randomly generated numbers from this distribution and rounded them to the closest integers as the new group sizes. For numbers smaller than one (larger than six), we forced them to be one (six). After the group size was decided, we calculated the number of non-treated individuals in each new group by multiplying the group size by the proportion of non-treated participants.

Next, for each non-treated participant, we simulated 20 trials of negative emotion ratings which was consistent with the experiment trial number. The simulated ratings in each proportion condition were generated from a distribution of ratings of the corresponding proportion in the experiment. This rating distribution was estimated first as a normal distribution based on the mean and standard deviation of ratings in each proportion condition. We then squeezed the range of the distribution to [1, 10] and rounded each number drawn from the distribution to the closest integer as the rating.

We kept generating new groups for every proportion condition until the total number of non-treated participants (both simulated and original numbers) reached a target number (243 participants per proportion condition). The target number was determined by the largest number of non-treated participants among all proportion conditions in the original experiment. An iteration was completed when the non-treated participant numbers in all proportion conditions were equal to or larger than the target number. Note that we excluded one proportion condition (83.3% treatment) from all simulation processes because it only had 4 non-treated participants in the original data. We were not able to estimate its distribution of ratings due to the very limited sample size. The result of each iteration was groups of 243–245 non-treated participants.

Text processing. We processed text in the task using AlephBert, which is a large pre-trained language model in Hebrew³⁸. Prior processing the text via AlephBert, we removed punctuation and capital letters from the text, replaced symbols with words, removed double spaces and made sure that there were spaces after commas and periods. We then used AlephBert to generate embeddings for each text that was produced by each participant in each trial. To generate the “pure reappraisal” vector we averaged all vectors of the all-reappraisal condition (all six participants in the reappraisal condition). Similar process was done to the non-treated condition: we averaged all the responses in the all non-treated condition (all six participants in the control condition). We then subtracted the non-treated ratings from the reappraisal, which produced a vector representation of “pure reappraisal”. We then compared the pure reappraisal vector to each of the participants’ responses using cosine similarity.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The data generated in all studies, including the pilot, main, and supplementary study, have been deposited on Open Science Framework⁵⁰ (<https://doi.org/10.17605/OSF.IO/DPCQ2>) and are publicly available.

Code availability

R Codes for main hypotheses testing and figure generation have been deposited under the OSF database (<https://doi.org/10.17605/OSF.IO/DPCQ2>).

References

1. Gross, J. J. Emotion regulation: Current status and future prospects. *Psychological Inq.* **26**, 1–26 (2015).
2. Gross, J. J. Antecedent- and response-focused emotion regulation: Divergent consequences for experience, expression, and physiology. *J. Personal. Soc. Psychol.* **74**, 224–237 (1998).
3. Wang, K. et al. A multi-country test of brief reappraisal interventions on emotions during the COVID-19 pandemic. *Nat. Hum. Behav.* **5**, 1089–1110 (2021).
4. Liu, J. J. W., Ein, N., Gervasio, J. & Vickers, K. The efficacy of stress reappraisal interventions on stress responsivity: A meta-analysis and systematic review of existing evidence. *PLoS One* **14**, e0212854 (2019).
5. Preuss, H., Capito, K., van Eickels, R. L., Zemp, M. & Kolar, D. R. Cognitive reappraisal and self-compassion as emotion regulation strategies for parents during COVID-19: An online randomized controlled trial. *Internet Interventions* **24**, 100388 (2021).
6. Bar-Tal, D., Halperin, E. & De Rivera, J. Collective emotions in conflict situations: Societal implications. *J. Soc. Issues* **63**, 441–460 (2007).
7. Halperin, E. *Emotions in Conflict: Inhibitors and Facilitators of Peace Making*. (Routledge, New York, NY, 2016).
8. Goldenberg, A., Halperin, E., van Zomeren, M. & Gross, J. J. The process model of group-based emotion: Integrating intergroup emotion and emotion regulation perspectives. *Personal. Soc. Psychol. Rev.* **20**, 118–141 (2016).
9. Davies, S., Pettersson, T. & Öberg, M. Organized violence 1989–2021 and drone warfare. *J. Peace Res.* **59**, 593–610 (2022).
10. Institute for Economics & Peace. *Economic Value of Peace 2021: Measuring the Global Economic Impact of Violence and Conflict*. <http://visionofhumanity.org/resources> (2021).
11. Wenzel, T., Schouler-Ocak, M. & Stompe, T. Editorial: Long Term Impact of War, Civil War and Persecution in Civilian Populations. *Front. Psychiatry* **12**, 733493 (2021).

12. Halperin, E., Porat, R., Tamir, M. & Gross, J. J. Can Emotion Regulation Change Political Attitudes in Intractable Conflicts? From the Laboratory to the Field. *Psychol. Sci.* **24**, 106–111 (2013).
13. Halperin, E., Pliskin, R., Saguy, T., Liberman, V. & Gross, J. J. Emotion regulation and the cultivation of political tolerance: Searching for a new track for intervention. *J. Confl. Resolut.* **58**, 1110–1138 (2013).
14. Hurtado-Parrado, C. et al. Emotion Regulation and Attitudes Toward Conflict in Colombia: Effects of Reappraisal Training on Negative Emotions and Support for Conciliatory and Aggressive Statements. *Front. Psychol.* **10**, (2019).
15. Porat, R., Erel, L., Pnueli, V. & Halperin, E. Developing ReApp: an emotion regulation mobile intervention for intergroup conflict. *Cognition Emot.* **34**, 1326–1342 (2020).
16. Manstead, A. & Fischer, A. H. Social appraisal: The social world as object of and influence on appraisal processes. in *Emotion: Theory, methods, research* (eds. Scherer, K. R., Schorr, A. & Johnstone, T.) 221–232 (Oxford University Press, New-York, NY, 2001).
17. Clément, F. & Dukes, D. Social appraisal and social referencing: Two components of affective social learning. *Emot. Rev.* **9**, 253–261 (2017).
18. Morris, A. S., Silk, J. S., Steinberg, L., Myers, S. S. & Robinson, L. R. The Role of the Family Context in the Development of Emotion Regulation. *Soc. Dev.* **16**, 361–388 (2007).
19. Silvers, J. A. Adolescence as a pivotal period for emotion regulation development. *Curr. Opin. Psychol.* **44**, 258–263 (2022).
20. Sorce, J. F., Emde, R. N., Campos, J. & Klinnert, M. D. Maternal emotional signaling: Its effect on the visual cliff behavior of 1-year-olds. *Developmental Psychol.* **21**, 195–200 (1985).
21. Hatfield, E., Cacioppo, J. T. & Rapson, R. L. *Emotional Contagion*. (Cambridge University Press, New York, NY, 1994).
22. Goldenberg, A. & Gross, J. J. Digital emotion contagion. *Trends in Cognitive Science* online (2020) <https://doi.org/10.1017/CBO9781107415324.004>.
23. Cialdini, R. B. & Goldstein, N. J. Social influence: Compliance and conformity. *Annu. Rev. Psychol.* **55**, 591–621 (2004).
24. Valente, T. W. Network Interventions. *Science* **337**, 49–53 (2012).
25. Cross, R. L. & Parker, A. *The Hidden Power of Social Networks: Understanding How Work Really Gets Done in Organizations*. (Harvard Business Press, 2004).
26. Paluck, E. L., Shepherd, H. & Aronow, P. M. Changing climates of conflict: A social network experiment in 56 schools. *Proc. Natl Acad. Sci.* **113**, 566–571 (2016).
27. Greetham, D. V., Sengupta, A., Hurling, R. & Wilkinson, J. Interventions in social networks: impact on mood and network dynamics. *Advs. Complex Syst.* **18**, 1550016 (2015).
28. Centola, D. The spread of behavior in an online social network experiment. *Science* **329**, 1194–1197 (2010).
29. Guilbeault, D. & Centola, D. Topological measures for identifying and predicting the spread of complex contagions. *Nat. Commun.* **12**, 4430 (2021).
30. Kim, D. A. et al. A Randomised Controlled Trial of Social Network Targeting to Maximise Population Behaviour Change. *Lancet* **386**, 145–153 (2015).
31. Cho, Y., Hwang, J. & Lee, D. Identification of effective opinion leaders in the diffusion of technological innovation: A social network approach. *Technol. Forecast. Soc. Change* **79**, 97–106 (2012).
32. Banerjee, A., Chandrasekhar, A. G., Duflo, E. & Jackson, M. O. The Diffusion of Microfinance. *Science* **341**, 1236498 (2013).
33. Grand, G., Blank, I. A., Pereira, F. & Fedorenko, E. Semantic projection recovers rich human knowledge of multiple object features from word embeddings. *Nat. Hum. Behav.* **6**, 975–987 (2022).
34. Pliskin, R., Halperin, E., Bar-Tal, D. & Sheppes, G. When Ideology Meets Conflict-Related Content: Influences on Emotion Generation and Regulation. *Emotion* **18**, 159–170 (2018).
35. Goldenberg, A., Saguy, T. & Halperin, E. How group-based emotions are shaped by collective emotions: Evidence for emotional transfer and emotional burden. *J. Personal. Soc. Psychol.* **107**, 581–596 (2014).
36. Goldenberg, A. et al. Amplification in the evaluation of multiple emotional expressions over time. *Nat. Hum. Behav.* **6**, 1408–1416 (2022).
37. Frijda, N. H. *The Emotions*. (Cambridge University Press, London, UK, 1986).
38. Seker, A. et al. AlephBERT: A Hebrew Large Pre-Trained Language Model to Start-off your Hebrew NLP Application With. Preprint at <https://doi.org/10.48550/arXiv.2104.04052> (2021).
39. Centola, D. & Macy, M. W. Complex contagions and the weakness of long ties. *Am. J. Sociol.* **113**, 702–734 (2007).
40. Goldenberg, A., Garcia, D., Halperin, E. & Gross, J. J. Collective emotions. *Curr. Directions Psychological Sci.* **29**, 154–160 (2020).
41. Yeager, D. S. et al. A national experiment reveals where a growth mindset improves achievement. *Nature* **573**, 364–369 (2019).
42. Walton, G. M. & Cohen, G. L. A Brief Social-Belonging Intervention Improves Academic and Health Outcomes of Minority Students. *Science* **331**, 1447–1451 (2011).
43. Kim, D. A. et al. Social network targeting to maximise population behaviour change: a cluster randomised controlled trial. *Lancet* **386**, 145–153 (2015).
44. Oveis, C., Gu, Y., Ocampo, J. M., Hangen, E. J. & Jamieson, J. P. Emotion regulation contagion: Stress reappraisal promotes challenge responses in teammates. *J. Exp. Psychol.: Gen.* **149**, 2187–2205 (2020).
45. Taylor, S. J., Muchnik, L., Kumar, M. & Aral, S. Identity effects in social media. *Nat. Hum. Behav.* **7**, 27–37 (2023).
46. Guilbeault, D., Becker, J. & Centola, D. Social learning and partisan bias in the interpretation of climate trends. *Proc. Natl Acad. Sci.* **115**, 9714–9719 (2018).
47. Pliskin, R., Bar-Tal, D., Sheppes, G. & Halperin, E. Are leftists more emotion-driven than rightists? The interactive influence of ideology and emotions on support for policies. *Personal. Soc. Psychol. Bull.* **40**, 1681–1697 (2014).
48. Wang, K. et al. A global test of brief reappraisal interventions on emotions during the COVID-19 pandemic. *Nat. Hum. Behav.* <https://doi.org/10.31234/osf.io/m4gppq>. (2020).
49. Gross, J. J. & John, O. P. Individual differences in two emotion regulation processes: Implications for affect, relationships, and well-being. *J. Personal. Soc. Psychol.* **85**, 348–362 (2003).
50. Pinus, M., Cao, Y., Coman, A., Gross, J. & Goldenberg, A. Data for Emotion Regulation Contagion. (2022) <https://doi.org/10.17605/OSF.IO/DPCQ2>.

Acknowledgements

A.G. acknowledges funding support by the National Science Association (Award number 1823410). The authors would like to thank Gil Ben Josef for his help with running the pilot study. We also would like to thank Almog Simchon for his consultation on the text analysis. Finally, we would like to thank Talma Goldenberg for her help in managing the timeline of the project.

Author contributions

M.P., Y.C., and A.G. designed research, conducted experiments, analyzed data, and wrote the manuscript. E.H., A.C., and J.J.G. reviewed and edited the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-025-56538-x>.

Correspondence and requests for materials should be addressed to Amit Goldenberg.

Peer review information *Nature Communications* thanks Douglas Guilbeault, Hana Shepherd and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025