Supplementary Information for

Social contagion in attitudes about prejudice

Table S1 contains the treatment effects for the full set of cohabitants and the set of cohabitants who were fewer than 26 years in age different from the directly contacted individual on the survey conducted three days after the intervention. Table S2 contains the treatment effects for the full set of cohabitants and the set of cohabitants who were fewer than 26 years in age different from the directly contacted individual on the survey conducted three weeks after the intervention. Table S3 contains the treatment effects for the full set of cohabitants and the set of cohabitants who were fewer than 26 years in age different from the directly contacted individual on the survey conducted six weeks after the intervention.

Figure S2 graphically displays the treatment effects across the three survey waves in standard deviations for the full set of cohabitants. This figure is analogous to Figure 1 from the main text, but uses the full set of cohabitants rather than only those who were relatively similar in age to the directly contacted individual.

Identifying cohabitants

I identify cohabitants using the data provided by Broockman and Kalla (2016) (henceforth, BK). In their data, BK identify 68,378 individuals who were mailed an invitation to participate in their survey. In households with two or three registered voters, all household members were invited to participate, with each individual receiving a separate code with which they could participate. In households with more than three registered voters, a random sample of three household members were invited to participate. Among those invited, 1,825 individuals enrolled via participation in the baseline survey, coming from 1,295 households. As such, an average of about 1.41 people per household completed the baseline survey. All households were then randomly assigned to either the treatment (n = 913) or control (n = 912) conditions. Using these random assignments, canvassers attempted to contact voters at their doors. Canvassers were successful in contacting 501 voters. These individuals constitute the directly contacted sample that BK originally analyzed. Among the original set of 1,825 survey respondents, 191 individuals cohabitate with an individual who was contacted directly, did not directly receive the treatment, and completed at least one follow-up survey. This set of 191 comprise the sample that I use to investigate contagion in the treatment effect.

Design assumptions

One of the key assumptions of the design is that the treatment and control groups, and by extension in this analysis the cohabitants of the treatment and control groups, do not differ pre-treatment. I tested this assumption by investigating pre-treatment covariates and the extent to which they differ based on treatment and control assignments. Tables S4, S5, S6, S7 contains the information on pre-treatment characteristics for the treatment and control groups for the full set of cohabitants and the set of cohabitants who were fewer than 26 years in age different from the directly contacted individual for the pre-treatment survey, the three-day survey, the three-week survey, the six-week survey, respectively. The only significant difference between the cohabitants of the treatment and control groups is that in the three-week and six-week surveys, black respondents were more likely to be in the control group than the treatment group.

Mediation of attitudes

One possible mediator for the change in attitudes in cohabitants is the attitude change of the directly contacted individual. I tested for this using the mediation testing procedures described by Tingley, Yamamoto, Hirose, Keele, and Imai (2014). I first estimated the same linear model as described in the main text, but including the directly contacted individual's attitude for the dependent variable used in the model in the same wave. That is, in the model for the omnibus index in the three-day surveys the model includes the directly contacted individual's measure for the omnibus index in the three-day surveys as a covariate. I then estimate a model in which the directly contacted individual's measure (the mediator) is the dependent variable. Then, I estimate the mediated effect, the direct effect, and the total effect, as presented in Table S8 (for the full sample of cohabitants) and in Table S9 (for the subsample of cohabitants who are 25 years in age or fewer different from the directly contacted individual). As the tables show, in neither of the models is there a significant mediation effect. This is somewhat surprising, as the most likely pathway for the effect of the intervention to reach a cohabitant is through the directly contacted individual. However, it is important to remember that the survey measure of the directly contacted individual's attitude may not be the best measure of the process through which the mediated relationship takes place. Perhaps a better mediator would be the amount of time the directly contacted individual spent talking with the cohabitant about the intervention. For example, it is possible that a directly contacted individual whose attitude was not changed would nonetheless pass on the information from the canvasser to a cohabitant through conversation and change that cohabitant's attitude. Future research may wish to better understand when and how a directly impacted individual passes on the effects of an intervention such as this one through the measurement of the nature of the relationship between connected individuals and the type of communication that takes place between them.

Using age difference as a moderator

The main text shows that among a subsample of the cohabitants (those who are 25 years in age different from the directly contacted individual) the treatment effect spills over. To test whether the difference in age moderates the relationship between treatment and a cohabitant's attitudes, I conducted regression analyses in which the treatment indicator is interacted with the absolute value of the difference in age between the directly contacted individual and the cohabitant. The results of these analyses are presented in Table S10 (for the continuous measurement of age difference) and S11 (for the dichotomized version of age difference, in which *close in age* measures relationships in which the directly contacted individual is 25 years in age different from the cohabitant or fewer). As the results show, there is little evidence that the treatment effect is moderated by age difference. Only one of the interaction tests shows statistical significance – the transgender tolerance index in the three-day surveys for both the continuous and dichotomized versions of the age difference measure. In this case, the interaction is in the theoretically expected direction: that as the difference in age increases, the impact of the treatment decreases. Most of the other interactions are in the expected direction, but they are too noisy to distinguish from chance. These results suggest that there is likely not a linear relationship between age difference and treatment effect. It is important to remember that age difference is a proxy for tie strength, and that in other studies (e.g., Bond et al., 2012; Jones et al., 2017) the likelihood of influence is related to tie strength in a non-linear fashion. Specifically, these studies find that the closest relationships are those that are most likely to show evidence of social influence.

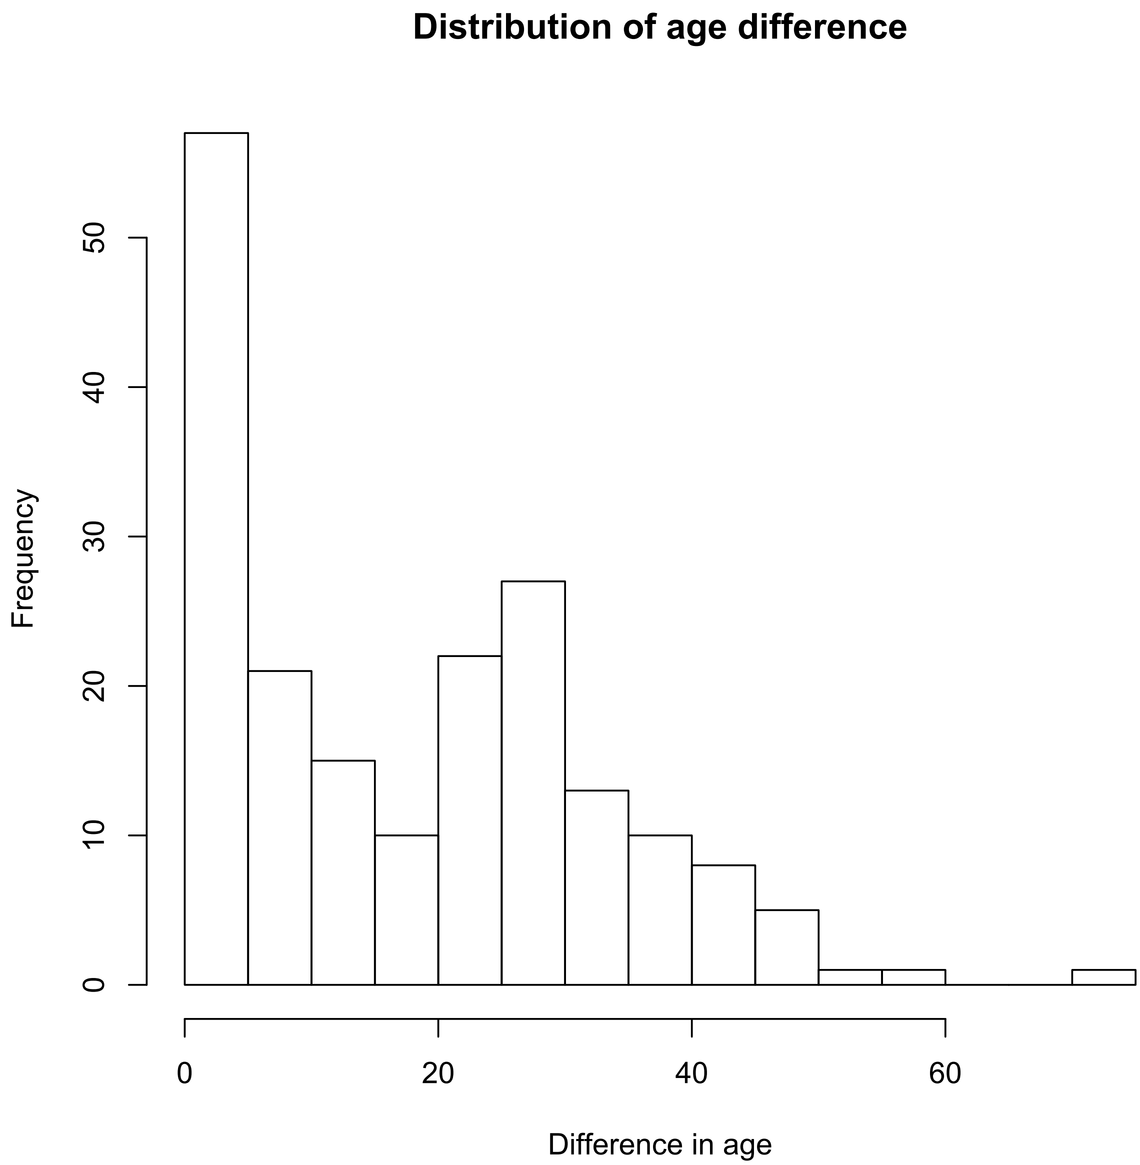

Figure S1. Distribution of age difference between directly contacted individual and cohabitant in years.

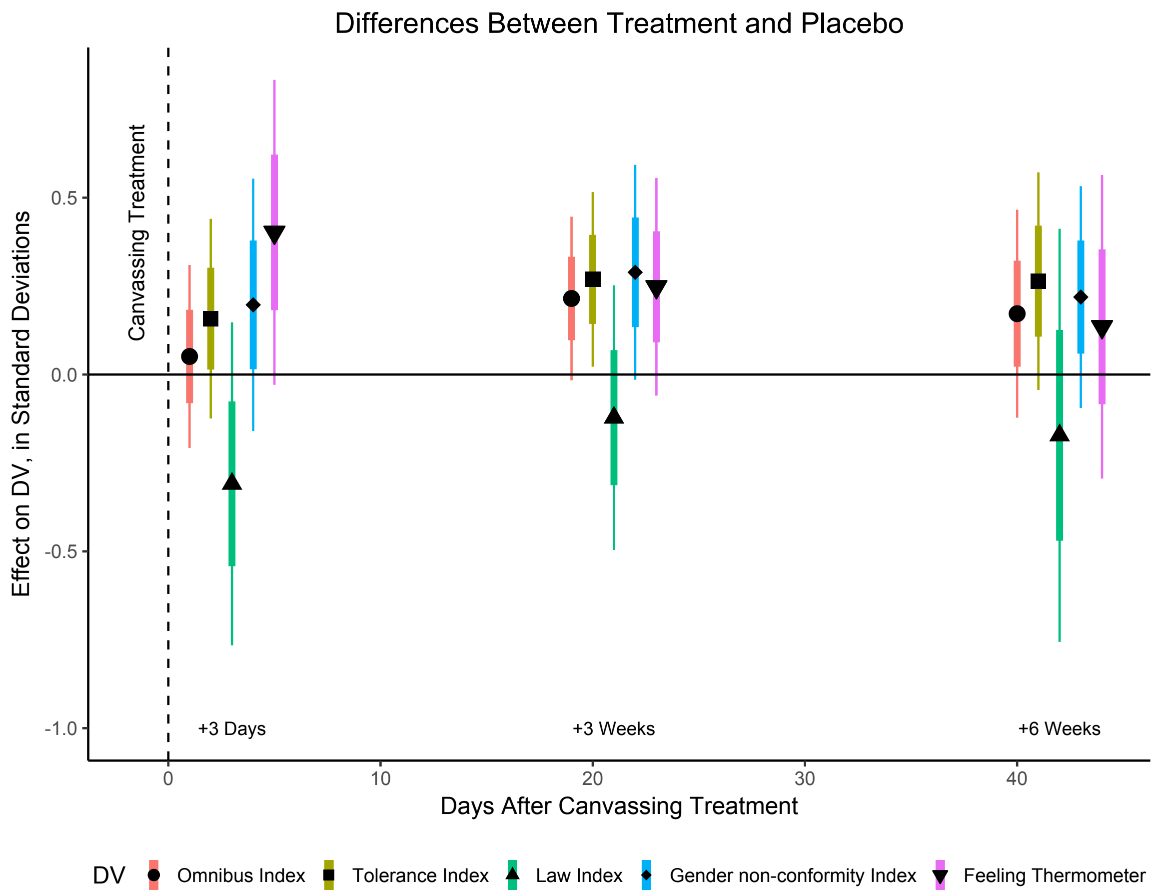

Figure S2. Complier average causal effects on all cohabitants of those exposed to the intervention. Points indicate estimated treatment effect of the intervention in standard deviations. Thicker lines are standard errors of the estimate and thinner lines are 95% confidence intervals.

References

Aral, S., & Walker, D. (2012). Identifying influential and susceptible members of social networks. *Science*. doi:10.1126/science.1215842

Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D. I., Marlow, C., Settle, J. E., & Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature, 489*(7415), 295-298. doi:10.1038/nature11421

Broockman, D. E., & Kalla, J. L. (2016). Durably reducing transphobia: A field experiment on door-to-door canvassing. *Science, 352*(6282), 220-224. doi:10.1126/science.aad9713

Jones, J. J., Bond, R. M., Bakshy, E., Eckles, D., & Fowler, J. H. (2017). Social influence and political mobilization: Further evidence from a randomized experiment in the 2012 US presidential election. *PLoS ONE, 12*(4). doi:10.1371/journal.pone.0173851

Tingley, D., Yamamoto, T., Hirose, K., Keele, L., & Imai, K. (2014). Mediation: R package for causal mediation analysis. *Journal of Statistical Software, 59*(5), 1-38. doi:10.18637/jss.v059.i05