



Research paper

Masks, cameras and social pressure[☆]Itzhak Rasooly^{a,b,*}, Roberto Rozzi^{c,*}^a Department of Economics, Sciences Po, 28 Rue des Saints-Pères, 75007 Paris, France^b Paris School of Economics, 48 Bd Jourdan, 75014 Paris, France^c Department of Economics and Statistics, Università di Siena, Piazza San Francesco 7, 53100 Siena, Italy

ARTICLE INFO

JEL classification:

D90

C93

C73

Keywords:

Social pressure

Field experiment

Conformity

ABSTRACT

In this paper, we report the results of two experiments that randomise the share of individuals who are taking an action in subjects' immediate environment. Despite the differences between our two settings (face masks and online camera use), we uncover some empirical results that are common to both. First, we find that the share of individuals taking the relevant action is increasing in the share of others who take the action (although the relationship need not be linear). Second, and despite this, we find that many individuals nonetheless defy social pressure. Our results point both to the importance of social pressure as well as its very real limits in our settings.

1. Introduction

There is a large literature demonstrating the power of peer effects and descriptive social norms across a range of domains. For instance, studies have found that we look to others when deciding whether to evade our taxes (Bott et al., 2020), donate to charity (Agerström et al., 2016) and even whether to vote (Gerber and Rogers, 2009). Of course, these examples are somewhat arbitrary: it is hard to think of even one activity that is not somehow shaped by our expectations about the behaviour of others.

Despite a multitude of experiments in this area, few studies evaluate the impact of exogenously varying behaviour in the immediate environment of subjects. Instead, they study the impact of shocking beliefs about prevalence in the wider population, typically through the transmission of statistical information. In part, this may be due to logistical reasons: it is much easier to randomise the provision of statistical information than to randomise the behaviour that subjects observe. In practice, however, social norms do not usually operate via statistical information; instead, individuals react to the behaviour of others that they observe around them.

To capture such social effects, we conduct experiments that directly randomise the share of individuals taking an action in subjects' immediate environment. Our experiments consider the impact of setting the share at many different values: we study how

[☆] For useful comments and discussions, we would like to thank Johannes Abeler, Loukas Balafoutas, Miguel Ballester, Michele Bernasconi, Ennio Bilancini, Adam Brzezinski, Valerio Capraro, Jean-Paul Carvalho, Luca Corazzini, Pietro Dindo, Andrea Gallice, Junnan He, Bernhard Kasberger, Vatsal Khandelwal, Eduard Krkoska, Marco LiCalzi, Valeria Maggian, Heinrich Nax, Paolo Pellizzari, Vivek Roy-Chowdhury, Karl Schlag, Jasmine Theilgaard, Séverine Toussaert, Marie Claire Villeval and Peyton Young. We would also like to thank audiences at the EEA-ESEM Congress, the CREST/CEFM Workshop on Experimental Economics, ASFEE Lyon, ESA Bologna, the University of Oxford, the University of Turin, and IMT Lucca. For excellent research assistance, we are grateful to Peter Brookes, Aniket Chakravorty, Mary Chen, Michelle Cho, Thomas Dingwall, Ethan Dodd, Juliet Dyrud, Jordan Edwards-Zinger, Leah Fahy, Kirsten Fletcher, Tim Green, Jacob Hill, Maisie Howard, Wade Kamphuis, Jasvin Khurana, Antonia Ledda, Amélie Marescaux, Ria McDonald, Declan Murray, Lucy Owen, Eunice Poon, Rose Poyser, Wyatt Radzin, Neily Raymond, Doaa Shabbir, Sahani de Silva, Nina Skrzypczak, Lunchen Song, Ivan Toth-Rohonyi, Natalie Vriend and Gwen Williams. Finally, we would like to thank Exeter College, the George Webb Medley Fund, and Ca' Foscari University of Venice, Italy for generous financial support.

Replication package.

* Corresponding authors.

E-mail addresses: itzhak.rasooly@sciencespo.fr (I. Rasooly), roberto.rozzi@unisi.it (R. Rozzi).

<https://doi.org/10.1016/j.jebo.2024.106699>

Received 9 October 2023; Received in revised form 31 July 2024; Accepted 5 August 2024

Available online 28 August 2024

0167-2681/© 2024 Elsevier B.V. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

prevalence varies depending on whether the share is 0%, 25%, 50%, 75%, or 100%. This gives us a rich picture of how subject behaviour varies in response to changes in the local prevalence of the behaviour.

Our first experiment concerns face mask usage. The basic idea of the experiment was straightforward. Subjects entered a room (one at a time) thinking that they were there solely to solve a decision problem involving lotteries. Unbeknownst to them, the number of the four experimenters in the room wearing a face mask had been randomised (leading to treatments in which 0/4, 1/4, 2/4, 3/4, or 4/4 experimenters were wearing a mask). We then observed whether each subject themselves chose to wear a face mask.

The experiment took place in Oxford over the course of nine days in February/March 2022. In total, we conducted fourteen three-hour sessions across twelve different colleges; and repeated our experimental protocol 646 times (each time with a different subject). Importantly, the experiment took place at a time in which face masks were no longer required by law or university rules, but still remained not abnormal. As a result, this was an ideal setting for capturing the implications of social pressure.

Our first experiment yielded three main results. First, we find that some individuals are susceptible to social pressure: the more experimenters who wear a mask, the more likely are subjects themselves to wear a mask. Second, and despite this, we find that a substantial share of individuals choose to wear a mask (or not wear one) no matter what others are doing. Third, we find that individual behaviour responds non-linearly to the share of experimenters wearing a mask. In particular, we observe an especially large jump when moving from the 3/4 to 4/4 treatments, consistent with an ‘everybody effect’ where social pressure becomes especially acute if everybody in the relevant environment chooses to do a particular activity. Notice that such non-linearity would not be detectable using a binary experiment, which illustrates one of the many advantages of our multi-treatment design.

In order to assess the robustness of our findings, we conducted an analogous experiment in a very different context: camera use in online calls. The idea of this experiment was also straightforward. Subjects joined a Zoom call (one at a time) knowing only that they were attending in order to participate in an economics experiment. Unbeknownst to them, the number of the four experimenters on the call with their laptop camera on had been randomised (leading again to five treatments, corresponding to 0/4, 1/4, 2/4, 3/4, and 4/4 experimenters with their camera on). We then observed whether each subject themselves chose to use their video camera. In total, we repeated this process 1114 times, leading to a sample size that was almost twice as large as that obtained in our first experiment.

Conducting this experiment led to similar, although not identical, results. We again find behavioural responses to social pressure: individuals are more likely to use a camera as the share of experimenters using a camera rises. We also again find high levels of non-compliance, with many participants choosing to use their cameras (or not) regardless of how many others are doing the same. Despite these similarities, the responsiveness of behaviour in this context is not precisely the same as that estimated in the mask setting; and appears to be substantially more linear.

Finally, we use our experimental results to calculate the distribution of individual thresholds across our two settings. The estimation procedure is straightforward: to calculate the share who ‘switch’ at a particular point, one simply compares choice frequencies across neighbouring treatment groups (e.g. the 0/4 vs. 1/4 treatments). In both experiments, we find that individual thresholds are very heterogeneous, in contrast to canonical models in evolutionary game theory (e.g. Young, 1993). We then use these thresholds to calibrate an evolutionary model in which agents adjust their actions in response to the actions of their neighbours. Our model predicts that, in the long run, individuals will take different actions despite engaging in copying-like behaviour — a point that can explain the common lack of behavioural conformity in settings where social pressure is nonetheless important.¹

Our results reveal both the importance of social pressure as well as its limits in real-world settings. For example, in the mask experiment, some individuals do switch to mask-wearing as the share of others wearing a mask rises — but a substantial fraction (51%) choose not to wear a mask even when all four others are doing so. This is somewhat surprising given that the experimenters were clearly visible, that masks were freely available, and that choosing to wear a mask was essentially costless in terms of time and effort. It is also striking given that the social pressure was induced not via statistical information (which might seem rather easier to ignore) but rather through direct exposure to the behaviour of others.

Our study contributes to a number of literatures across economics and related disciplines. First, our study builds on the conformity literature following Asch (1951). In contrast to this literature, our study concerns individuals’ actions (e.g. whether to wear a mask) as opposed to their cognitive judgements. Perhaps more importantly, our study also uses randomisation across many different treatments, in contrast to the ‘binary’ experiments in the Asch paradigm (see Bond and Smith, 1996 for an overview). Despite these differences, our study reinforces (Asch, 1951)’s findings on both the importance and limits of social pressure (see Friend et al., 1990 for discussion). The non-conformity that we document is arguably more striking: although it is perhaps unsurprising that some individuals stick to their judgements when facing a simple cognitive task, it is more surprising that a large fraction stick to their preferred decision even in settings where there is no ‘correct’ action and conforming (e.g. by wearing a mask) is relatively costless.²

Second, our study contributes to the broader literature on the importance of peer effects and descriptive social norms. The current literature consists of a series of generally binary experiments across a variety of domains (see Cialdini, 2007; Mascagni, 2018; Farrow

¹ This result contrasts with the standard predictions of models in evolutionary game theory. In such models (see, e.g. Kandori et al., 1993; Young, 1993; Jackson and Yariv, 2007; Young, 2009; Kreindler and Young, 2013) agents are identical and thus have a common ‘threshold’; this means that the system converges to a state in which either everybody or nobody does the activity. In contrast, our models converge to a ‘mixed’ equilibrium with heterogeneous choices. See also Centola and Baronchelli (2015), Centola et al. (2018), Andreoni et al. (2021) and Ehret et al. (2022) for experimental work on the dynamics of systems in which agents have conformity incentives.

² In part, this might be because the subjects had experience with the relevant activities (e.g. mask wearing) prior to the experiment and thus had the opportunity to build up habits along with beliefs about the desirability of the relevant action.

et al., 2017 for reviews).³ In contrast, our study is the first to randomly confront subjects with many different levels of prevalence in their immediate environment; and the first to do so in any setting (not just the settings of face masks and video calls).

Third, and more narrowly, we contribute to the literature on the social determinants of face mask wearing. The existing papers in this literature rely either on vignette-based experiments and surveys (Bokemper et al., 2021; Barceló and Sheen, 2020; Rudert and Janke, 2021; Goldberg et al., 2020; Barile et al., 2021) or instead on observational data (Freidin et al., 2022; Woodcock and Schultz, 2021). We contribute to this literature by conducting the first ever randomised field experiment on the social determinants of face mask use. Our use of randomisation allows us to side-step some of the issues that afflict the previous studies in this area.⁴

Fourth, we contribute to the literature on the social determinants of video camera use. Existing papers in this literature are again based on surveys: see, for example, Castelli and Sarvary (2021), Gherheş et al. (2021), Sederevičiūtė-Pačiauskienė et al. (2022) and Bedenlier et al. (2021). Our study is the first to examine this topic through use of a randomised field experiment; again, this allows us to avoid various econometric issues that would otherwise arise in the causal estimation of social effects.

The remainder of this article is structured as follows. Section 2 outlines the design of our face mask experiment and the associated results. Section 3 presents the design and results for our experiment on video cameras. Section 4 discusses what these results might mean for the distribution of individual thresholds and long-run behaviour. Finally, Section 5 concludes with a discussion of future research suggested by our experiments.

2. Masks

2.1. Experimental design

We now describe our first experiment aimed at quantifying individual responsiveness to social pressure. The basic idea of the experiment was straightforward. Subjects entered a room thinking that they were there solely to solve a decision problem involving lotteries. Unbeknownst to them, the number of experimenters in the room wearing a face mask had been randomised. We then observed whether each subject themselves chose to wear a face mask (and how this varied with the number of experimenters wearing a mask in their immediate environment).⁵

This first experiment took place in Oxford in late February and early March of 2022. At this time, masks were not required by either law or university rules – however, they were also not unusual. These facts allowed us to avoid triviality: if masks were required by law (or a very strong social norm), then subjects would presumably wear a mask regardless of others' behaviour; and if masks were extremely unusual, it is unlikely that subjects could be induced to wear a mask by the mask wearing of others. In total, we conducted 14 three-hour sessions in 12 different colleges over 7 days (with the help of 16 research assistants, some of whom participated in multiple sessions). On average, around 46 participants attended each session; which led to a total sample size of 646 experimental subjects (see Table A.1 for the distribution of subjects across treatment groups).⁶

The structure of the experiment was as follows:

1. Each subject was asked to arrive at a room at a particular time.
2. Before each subject entered the room, the number of the four experimenters in the room who were wearing a mask (and the allocation of masks to experimenters) had been randomised. Thus, there were five treatment groups, corresponding to: 0/4 masks, 1/4 masks, 2/4 masks, 3/4 masks, 4/4 masks. We denote these treatments by T0, T1, etc.
3. Once a subject entered, they were asked to sit at a table in a way that gave them a clear view of the four experimenters. On the table were a box of masks as well as a bottle of hand sanitiser (such a set-up was common within the University of Oxford at the time). As a result, any subject who wished to wear a mask was able to do so.
4. Once the subject had sat down, each of the four experimenters introduced themselves by stating their name and subject of study. The purpose of this was to further ensure that each subject fully processed the number of experimenters who were wearing a mask.
5. The subject was asked their name, age, college and subject of study; and then given a decision problem involving lotteries.

³ Studies which study social norms and social pressure in various contexts include: Cialdini et al. (1990), Cason and Mui (1998), Ichino and Maggi (2000), Borsari and Carey (2003), Heldt (2005), Fortin et al. (2007), Goldstein et al. (2008), Martin and Randal (2008), Krupka and Weber (2009), Gerber and Rogers (2009), Allcott (2011), Ferraro and Price (2013), Ayres et al. (2013), Costa and Kahn (2013), Bursztyn et al. (2014), Damm and Dustmann (2014), Smith et al. (2015), Thöni and Gächter (2015), Efferson et al. (2015), Lefebvre et al. (2015), Allcott and Kessler (2019), Novak (2020), Kessler et al. (2021) and Linek and Traxler (2021). The most relevant of these are Krupka and Weber (2009) and Kessler et al. (2021), who vary information about past behaviour in the respective contexts of a dictator game and public good game. In contrast to these studies, we conduct field (not laboratory) experiments that study real-world behaviours, use our results to estimate the distribution of thresholds, and obtain much larger samples.

⁴ For example, attempts to study this problem using hypothetical questions (as in Bokemper et al., 2021) suffer from the issue that individuals may not know what they would do in a hypothetical situation — an especially pressing concern since imitative behaviour may well rest on unconscious cognition. Meanwhile, attempts to study this problem using observational data (as in Woodcock and Schultz, 2021) can suffer from both omitted variable bias and reverse causality issues (see Manski, 1993 for an influential exposition of this latter point). Our randomised experiment avoids both of these issues.

⁵ The experiment received approval from the University of Oxford's Departmental Research Ethics Committee (ECONCIA21-22-50). In line with the recommendations of the committee, we told subjects in advance that taking part in the experiment might involve interacting with unmasked individuals (which was common at the University of Oxford at the time). We also took reasonable social distancing precautions, including making sure that the experimental settings were well ventilated. We should also emphasise that, although we did not reveal the main purpose of our experiment to participants (as is not unusual in social science experiments), we did not explicitly deceive participants at any stage.

⁶ The experimental design (without an analysis plan) was pre-registered here: <https://www.socialscisearch.org/trials/9013>.

6. We then asked the subject to leave, and repeated the process for the next subject (see Appendix C for a more detailed description of the experimental protocol which includes the decision problem).

We recorded whether each subject was wearing a mask when they entered the room (this variable is labelled ‘pre’ in the tables). We also recorded whether they chose to wear a mask while interacting with the experimenters. More precisely, subjects were recorded as having decided to wear a mask if they put a mask on at any stage during the interaction (or instead entered wearing a mask and kept it on throughout).⁷ Finally, we recorded their choice in the lottery problem; as well as whether they asked if they ought to wear a mask (in such cases, each was told ‘it’s up to you’ by the data recorder).

Based on post-experimental conversations, it seemed that most subjects believed that our goal was to measure risk aversion. Importantly, none of the subjects appeared to suspect that the experiment had anything to do with face masks; and there was nothing in the experimental design that could have revealed this.⁸ This is reassuring since subjects might have acted in unnatural and unrepresentative ways if they had known that they were taking part in a face mask experiment.

Once all experimental sessions had been completed, we debriefed all subjects on the underlying purpose of the experiment. During the debriefing, subjects were given the opportunity to take part in an online survey. In the survey, subjects were asked to imagine that they entered a room and saw 4 people sitting around a table. They were then asked if they would wear a mask if none of the 4 people were wearing a face mask, if 1 of the 4 people were wearing a face mask, and so forth. Finally, they were asked to give an explanation for their answers, as well as whether they had contracted COVID-19 at any point during the pandemic. The purpose of the follow-up survey was to obtain some suggestive evidence on mechanisms, as well as some data on individual level decision rules (see Section 4 for discussion).

Before proceeding, we briefly discuss some issues raised by the experimental design. First, we should emphasise that our design picks up responses to changes in subjects’ immediate environment (namely, the closed room in which the experiment was performed). This environment is likely to be relevant since the people around a subject are exactly the people who might judge a subject for wearing (or not wearing) a mask. Nonetheless, one could imagine that subjects could also conceivably be influenced by the mask wearing of individuals who are not in their immediate environment (as well as various non-social factors, e.g. whether the subject enjoys wearing a mask). Our experiment was not designed to study the importance of these factors.

Second, we should emphasise that the exact results generated by our experimental design may not generalise to other tasks and cultural contexts (Henrich et al., 2010). However, one might hope that the more qualitative insights do generalise. We provide some evidence on this in Section 3, which implements the same experimental design using a different task. Additional experiments that study whether our qualitative results do extend across tasks and cultures would seem to be a useful area for future research.

2.2. Results

We now turn to our main results, beginning with a brief description of our sample. As shown by Table A.2, our average participant was around 21 years old; and approximately half of our sample was male. Participants were fairly evenly distributed across subject divisions, with social science students being most represented (33% of the sample). Turning to Table 1, we see that genders, subjects and ages were reasonably balanced across our five treatment groups. However, we do observe some imbalance in the share of participants who entered the room wearing a mask: for example, the share is 27% in treatment T2 but only 14% in T0. Given that this variable turns out to be highly predictive for our outcome (whether participants chose to wear a mask), we control for it in our main specification.

Our regressions take the form

$$y_i = \beta_0 + \sum_{j=1}^4 \beta_j T_i^j + \gamma x_i + u_i, \quad (1)$$

where y_i tracks whether individual i chose to wear a mask, T_i^j tracks whether i was assigned to treatment j , x_i is a vector of covariates (including whether i entered the room wearing a mask), and u_i is an error term. As usual, $(\beta_0, \beta_1, \dots, \gamma)$ are the parameters to be estimated; note that β_j is the expected difference in mask wearing rates between treatments T_j and T_0 . In our main specification, we control for participant age, gender, and whether they entered the room wearing a mask (the ‘pre’ variable). However, we also report uncontrolled regressions, as well as regressions that use the full set of controls that are available (including session and college fixed effects).

Fig. 1 plots the results from our main specification (see Table 2 for the corresponding estimates, and Tables A.3 and A.4 for the near identical results obtained by estimating probit and logit regressions). The x -axis indicates the treatments, expressed as the fraction of experimenters wearing a mask (0, 0.25, 0.5, 0.75, 1). The y -axis displays the predicted share of individuals wearing a mask in each treatment. To obtain this predicted share, we set the three control variables (age, gender, and pre) equal to their mean values; so we are implicitly correcting for any imbalance in the pre variable.

Several features of the data are apparent. First, we find evidence that the frequency of mask wearing is (weakly) monotonically increasing in the share of experimenters who wear a mask. This pattern is evident in all the specifications that we estimate, regardless

⁷ We did not observe any subjects who first put on a mask and then took it off.

⁸ We also required all research assistants to sign an agreement specifying that they would keep the main purpose of the experiment confidential throughout its duration.

Table 1
Balance table (experiment 1).

Variable	T0	T1	T2	T3	T4	<i>p</i> -value
Age	21.0 [.361]	21.3 [.539]	20.1 [.165]	20.6 [.219]	20.8 [.268]	.143
Pre	.142 [.031]	.157 [.032]	.266 [.039]	.242 [.039]	.203 [.035]	.060
Male	.535 [.044]	.522 [.043]	.461 [.044]	.548 [.045]	.421 [.043]	.189
Humanities	.323 [.042]	.246 [.037]	.250 [.038]	.347 [.043]	.256 [.038]	.237
Social	.268 [.039]	.403 [.043]	.336 [.042]	.298 [.041]	.353 [.042]	.177
MPLS	.213 [.036]	.209 [.035]	.305 [.041]	.242 [.039]	.233 [.037]	.380
Medical	.181 [.034]	.104 [.027]	.102 [.027]	.105 [.028]	.143 [.030]	.235

Notes. This table shows the average value of various variables across the five treatments. The variables are age, whether the subject entered wearing a mask ('pre'), gender, division of study (Humanities; Social Sciences; Mathematical, Physical & Life Sciences; Medical Sciences). The final column reports the *p*-value obtained from regressing the relevant variable on all treatment dummies and testing the hypothesis that the coefficients on all treatment dummies are equal to zero.

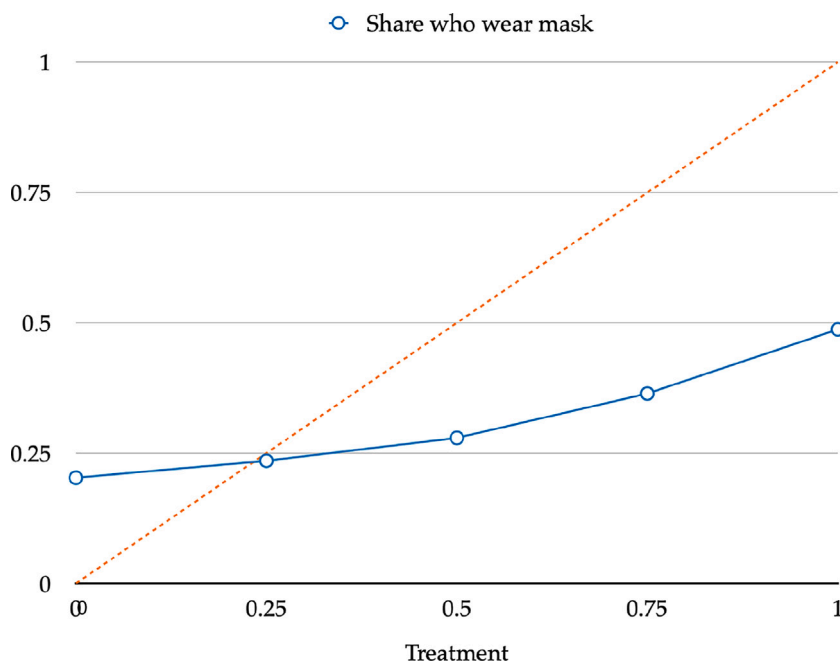


Fig. 1. Mask wearing by treatment group.

of whether they include controls, use logit or probit, etc. (again, see Tables 2, A.3 and A.4). From a statistical point of view, we can reject the hypothesis that treatment k leads to the same levels of mask wearing as treatment k' for all $k > k'$ with the exception of T1/T0 and T2/T1 (see Table A.5). While we discuss mechanisms later on, we note that this is consistent with a model in which higher rates of mask wearing lead to greater social pressure to wear a mask.⁹

As shown by Table A.6, most of the monotonicity comes from those who enter the room without wearing a mask. Effects become somewhat stronger once we restrict attention to these individuals; and the effects almost disappear entirely once we drop them from the sample. One obvious explanation for this asymmetry is that those who enter wearing a mask, despite the generally low prevalence of mask-wearing at the time, are generally very committed to wearing a mask no matter what. It should not be surprising that the decisions of these types are broadly unaffected by the decisions of those around them.

⁹ As stressed by Bicchieri (2016), these effects could be driven not by rational deliberation but rather by automatic and unconscious responses to cues. While this distinction may not matter for the behavioural implications of our results (see, e.g., Schelling, 1971 and Granovetter, 1978), it would be interesting to study in future work.

Table 2
Regressions (experiment 1).

Variable	No controls	Main specification	All controls
Treatment 1	.044 [.048]	.032 [.029]	.020 [.033]
Treatment 2	.171*** [.053]	.078** [.032]	.075** [.035]
Treatment 3	.238*** [.055]	.163*** [.039]	.156*** [.041]
Treatment 4	.331*** [.054]	.284*** [.043]	.289*** [.046]
Pre		.757*** [.029]	.741*** [.035]
Age		.002 [.005]	.001 [.005]
Male		-.007 [.026]	-.007 [.028]
Constant	.157*** [.032]	.014 [.107]	.130 [.144]
<i>n</i>	646	646	646
<i>R</i> ²	0.070	0.494	0.517

Notes. This table reports our main regressions. To obtain the estimates in the first column, we regress whether subjects wore a mask on the treatment dummies. In the second column, we control for subject age, gender, and whether they entered wearing a mask. The third column also includes session and college fixed effects. Robust standard errors in parentheses (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Second, we see that many individuals defy social pressure. In the treatment in which no experimenters wear a mask (T0), 20.0% of the participants nonetheless choose to wear a mask, a share which is statistically different from zero ($p < 0.0001$).¹⁰ In the language of Angrist et al. (1996), these people can be interpreted as ‘always wearers’, i.e. individuals who choose to wear a mask no matter how many others are doing the same (see Section 4 for elaboration). Similarly, in the treatment in which all experimenters wear a mask (T4), only 48.7% choose to wear a mask, which is again statistically different from 1 ($p < 0.0001$). The remaining 51.3% of individuals (who do not wear a mask) can be interpreted as ‘never wearers’, i.e. individuals who will never choose to wear a mask, no matter how many others are doing so (again, see Section 4 for a more formal discussion of this point).

Similar results are available if we look at changes. In treatment T0, out of the participants who entered the room wearing a mask, only 5.6% chose to take off their mask (see Table A.7). Similarly, in treatment T4, out of the participants who entered the room not wearing a mask, only 36.8% chose to put on a mask. It is quite striking that the majority of those who entered without a mask in T4 decided to defy social pressure in this way, especially given that all four experimenters were clearly visible and that a box of masks was available.

Third, the effect of the treatments on behaviour appears to be non-linear. Estimating a model with a quadratic term suggests some convexity ($p = 0.04$): see Table A.8.¹¹ Insofar as estimates appear non-linear, this is due to a large jump between the 3 and 4 treatments (the difference is 12 percentage points, in contrast to the average difference between treatments of 7 percentage points). This is indicative of a potential ‘everybody effect’, i.e. that a particularly large change in behaviour is induced by changing the share who are doing an action from ‘most people’ to ‘everybody’.

Taken together, our results reinforce findings in the Asch (1951) paradigm on both the power and limits of social pressure. In typical (Asch, 1951) conformity experiments, individuals are asked a question with an obviously correct answer. In contrast, in our experiment, subjects must make a real-world decision that cannot be ‘correct’ or ‘incorrect’ (as discussed earlier, our experiment also departs from Asch (1951) by conducting multi-treatment randomisation). For this reason, the non-conformity that we document is rather more striking: while it is unsurprising that some individuals cannot be convinced of something obviously false by social pressure, it is more surprising that many individuals refuse to change their actions in response to social pressure. To some extent, the non-conformity that we document might be due to individuals’ past experiences with mask wearing — although exploring this more fully is outside of the scope of the present paper.

Before moving to our second experiment, we briefly discuss the results of our online follow-up survey ($n = 120$).¹² As explained earlier, this survey directly asked participants how their decision to wear a face mask would vary with the number of individuals in the room who were also wearing a face mask. Given that individuals might not always know what they would do in a hypothetical situation, we do not emphasise the estimated decision rules obtained from this survey (although, reassuringly, they are also monotone

¹⁰ This p -value corresponds to a t -test of the null hypothesis that the population share equals zero; corresponding remarks apply to subsequent tests that share who always or never take an action equals a particular value.

¹¹ To further investigate whether adding a quadratic term allows us to better fit the data, we also use Lasso to select the most predictive variables out of: masks (the number of experimenters wearing a mask), masks², age, gender, pre, and a constant. The preferred model includes just four variables: masks, masks², pre, and the constant. Thus, while automatic model selection drops the age and gender controls, it retains the quadratic term.

¹² Although all subjects were invited to the follow-up survey, only 120 subjects completed it. Given the possibility of sample selection, one cannot automatically assume that results for the follow-up survey are entirely representative of results for the entire sample.

increasing in the number of mask wearers). However, we use the follow-up survey to address two issues that our original experiment could not speak to, namely individual level decision rules and mechanisms.

Our first finding from the online survey is that individual decision rules are plausibly monotone in the share of experimenters who are wearing a face mask. Indeed, over 99% of subjects report weakly increasing decision rules: if such subjects chose to wear a mask in some treatment T_k , they would also choose to wear the mask in all treatments $T_{k'}$ for $k' > k$. This finding helps validate our assumption in Section 4 that individual preferences have a threshold representation, which in turn provides an insightful decomposition of the observed aggregate behaviour. We should perhaps also stress that this finding cannot be obtained from the data from our main experiment, which is in principle consistent with the possibility that many individuals have decreasing decision rules.

Second, we obtain some suggestive evidence on why individuals are more likely to wear a mask if they see more mask wearing in their immediate environment. To do this, we consider only those individuals who reported that they would change their mask-wearing behaviour depending on the share of others wearing a mask. We then placed the explanations given by subjects into various categories, including whether they were trying to avoid being judged, trying to put others at ease, or taking high rates of mask wearing as a sign of high COVID risk levels (see Appendix D for a more detailed explanation of our categories along with examples). The main message from this exercise is that the health-based mechanism (i.e. that masks are used as a signal of COVID rates) is extremely unlikely to be driving our results: see Table A.9 for details. Instead, the observed changes seem to be driven by a variety of social learning and social pressure mechanisms, although exactly identifying the relative importance of these mechanisms is challenging.¹³

3. Cameras

3.1. Experimental design

In order to study the generality of our results, we conducted a second experiment which used a near-identical methodology in a very different context. The basic idea of this second experiment was the following. Subjects joined a Zoom call knowing solely that they were taking part in some kind of economics experiment. Unbeknownst to them, the number of experimenters on the call with their video camera on had been randomised. We then observed whether each subject themselves chose to use their camera. Thus, this second experiment was essentially the same as the first, except with the subject of video-camera instead of face mask usage.¹⁴

This second experiment took place online in late July and early August of 2022. We conducted 16 two-hour sessions over the course of 8 days (with the help of 20 research assistants, some of whom participated in multiple sessions). On average, each session was attended by around 70 participants, leading to a sample size of 1113 participants in total (see Table A.10 for the distribution of subjects across treatment groups). We recruited all participants from Prolific and restricted the sample to UK residents. All participants were required to have a working microphone and video camera.¹⁵

The structure of the experiment was as follows:

1. Each subject was asked to join a Zoom call at a particular time.
2. Before each subject joined the call, the number of the four experimenters in the meeting with their camera on (and which experimenters had their camera on) had been randomised. Thus, there were again five treatment groups: 0/4 cameras (denoted treatment T0), 1/4 cameras (T1), 2/4 cameras (T2), 3/4 cameras (T3), 4/4 cameras (T4).
3. Once a subject joined the call, all four experimenters introduced themselves by stating their name. The purpose of this was to ensure that each subject fully processed the number of experimenters whose cameras were on.
4. The subject was asked for their age, and whether they would want to donate half of a hypothetical £10 bonus payment to the next subject on the call.
5. We then asked the subject to leave the call, and repeated the process for the next subject (again, see Appendix C for a more detailed description of the experimental protocol).

Similarly to before, we recorded whether each subject had already turned their camera on when they joined the call; and whether they chose to use their camera at any point while interacting with the experimenters. We also recorded their choice in the decision problem; as well as whether they asked if they ought to turn their camera on (in such cases, each was told that 'it's up to you'). Finally, if a subject had not turned their camera on at any point during the call, we asked them if there were any issues with their video camera.¹⁶

¹³ One particular issue is that individuals may not be entirely honest about the reasons for their behaviour. For example, they might overstate the extent to which their behaviour is driven by altruistic reasons (e.g. trying to put others at ease), as opposed to a fear of being judged.

¹⁴ This experiment also received approval from the University of Oxford's Departmental Research Ethics Committee (ECONCIA21-22-44).

¹⁵ The experimental design (without an analysis plan) was pre-registered here: <https://www.socialscicenter.org/trials/9829>.

¹⁶ Unsurprisingly, asking this question occasionally had the effect of prompting participants to turn their video camera on. In such cases, we still recorded such participants as having chosen to not use their camera (on the basis that they had chosen not to use their camera until effectively asked to do so).

3.2. Results

We now turn to our results, beginning again with a description of our sample. In contrast with the student population studied in our first experiment, the average participant in this experiment was around 42 years old, with a standard deviation of 13.9 years (see Table A.11). Around 46% of the sample was male. As shown by Table 3, ages and genders were reasonably balanced across each of our five treatments. However, we again observe some imbalance in the share who joined the call with their camera on (the ‘pre’ variable), and so control for this variable in our main specification.

Table 3

Balance table (experiment 2).

Variable	T0	T1	T2	T3	T4	<i>p</i> -value
Age	42.2 [.940]	43.4 [.931]	42.3 [.903]	41.3 [.906]	42.7 [.990]	.615
Pre	.116 [.021]	.039 [.014]	.058 [.016]	.074 [.017]	.070 [.018]	.039
Male	.472 [.033]	.441 [.035]	.439 [.033]	.455 [.032]	.516 [.034]	.486

Notes. This table shows the average value of various variables across the five treatments. The variables are age, whether the subject joined the call with their camera on (‘pre’), and gender. The final column reports the *p*-value obtained from regressing the relevant variable on all treatment dummies and testing the hypothesis that the coefficients on all treatment dummies are equal to zero.

Our regressions take the same form as Eq. (1). That is, we regress whether an individual used their camera on the treatment dummies (using treatment T0 as the omitted category), and a vector of covariates. In our main specification, we control for participant age, gender, and whether they joined the call with their camera on. However, we once again also report uncontrolled regressions, as well as regressions that include the full set of possible controls (including session fixed effects).




Fig. 2. Camera use by treatment group.

Fig. 2 plots the results from our main specification (see Table 4 for the corresponding estimates, and Tables A.12 and A.13 for the near identical results obtained by estimating probit and logit regressions). Several points are apparent. First, similarly to the face mask experiment, we once again observe a monotone response to social pressure: the frequency of camera use is increasing in the share of experimenters who use a camera. This pattern arises in all of the specifications we estimate (see Tables 4, A.12 and A.13). In our main specification, we can reject the hypothesis that treatment k and treatment $k + 1$ lead to the same rates of camera usage ($p < 0.05$) for all k except $k = 3$; and we can always reject the hypothesis that treatment k and treatment $i + 2$ lead to the

Table 4
Regressions (experiment 2).

Variable	No controls	Main specification	All controls
Treatment 1	.077* [.043]	.118*** [.040]	.125*** [.041]
Treatment 2	.176*** [.043]	.209*** [.039]	.214*** [.044]
Treatment 3	.281*** [.043]	.308*** [.039]	.320*** [.049]
Treatment 4	.355*** [.044]	.380*** [.041]	.386*** [.057]
Pre		.579*** [.033]	.581*** [.034]
Age		.000 [.001]	.000 [.001]
Male		.024 [.027]	.023 [.027]
Constant	.241*** [.028]	.155*** [.047]	.094 [.061]
<i>n</i>	1113	1111	1109
<i>R</i> ²	0.069	0.161	0.183

Notes. This table reports our main regressions. To obtain the estimates in the first column, we regress whether subjects used a camera on the treatment dummies. In the second column, we control for subject age, gender, and whether they joined the call with their camera on. The third column also includes session fixed effects. Robust standard errors in parentheses (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

same rates of camera usage ($p < 0.01$) — see [Table A.14](#) for details. As before, this monotonicity is consistent with a model in which higher rates of camera use lead to greater social pressure to use a camera.¹⁷

Second, we once again observe that many individuals defy social pressure. In the treatment in which no experimenters use a camera (T0), 20.9% of the participants nonetheless choose to use a camera, a share which is statistically different from zero ($p < 0.0001$). As explained in [Section 4](#), such participants can be interpreted as ‘always users’, i.e. individuals who use a camera no matter how many others do the same. Similarly, in the treatment in which all experimenters use a camera (T4), only 58.7% choose to use a camera, which is again statistically different from 1 ($p < 0.0001$). The remaining 41.3% of individuals (who do not use a camera) can be interpreted as ‘never users’, i.e. individuals who will never choose to use a camera, no matter how many others are doing so. As before, similar results can be obtained by examining changes — see [Table A.16](#).

Third, the effect of the treatments in this context appears to be more linear. Statistically, we cannot reject a linear model: see [Table A.17](#). However, the jump between the 0 and 1 treatments (about 12 percentage points, in the main specification) is larger than the other 3 jumps (which are 9 percentage points, 10 percentage points, and 8 percentage points respectively). This provides some suggestive evidence on non-linearity, although one would need to obtain a larger sample to investigate this issue in greater detail.

4. Discussion

The previous sections show how individual behaviour varies with the prevalence of the behaviour amongst others in their immediate environment. In this section, we discuss *why* we observe the results that we do. We also discuss what our results might suggest for long run behaviour once used to calibrate an evolutionary model.

Explaining the results. To provide a formal explanation for our results, let us assume that individual preferences have a threshold representation. That is, we assume that, for every individual, there exists a threshold such that the individual does the action if and only if the number of others in the room who do the action exceeds this threshold. Such an assumption seems very plausible in our two contexts; and it is given some empirical validation by the online survey discussed in [Section 2.2](#).

Under this assumption, one can estimate the distribution of individual thresholds.¹⁸ To see how this works in practice, consider the data from the face mask experiment and define p_i as the share with a threshold of i , for $i \in \{0, 1, 2, 3, 4\}$. That is, for such values of i , p_i is the share of individuals who take the action if and only if they observe i or more of the four people in the room doing the same. Let us also define p_5 as the share who *never* do the action; and observe that $p_5 = 1 - \sum_{i=0}^4 p_i$.

[Table 5](#) reveals how the expected frequency of mask-wearing depends on the p_i parameters. In treatment T0, the only type who will do the action are the ‘always doers’, so the predicted share is p_0 . In treatment T1, the types who do the action are the ‘always doers’ in combination with those who switch when they see one person doing the action. More generally, in treatment Tk, the expected share who will do the action is $\sum_{i=0}^k p_i$. Given this, one can estimate the p_i by matching the parameters with the sample frequencies (as suggested, e.g., by maximum likelihood). For example, we obtain the estimates $\hat{p}_0 = 0.203$; and obtain $\hat{p}_1, \hat{p}_2, \hat{p}_3$,

¹⁷ As shown by [Table A.15](#), we once again see that the monotonicity is largely generated by those with the ‘pre’ variable equal to zero: see previous discussion.

¹⁸ Since we have one observation per subject, it is worth noting that we cannot estimate each subject’s threshold (but just the distribution).

\hat{p}_4 by computing the difference of mask wearing between neighbouring treatments. Finally, our estimate for the ‘never doers’ is obtained using $\hat{p}_5 = 1 - \sum_{i=0}^4 \hat{p}_i$.¹⁹

Table 5
Thresholds (experiment 1).

Treatment	Frequency	Predicted frequency
0	0.202	p_0
1	0.234	$p_0 + p_1$
2	0.280	$p_0 + p_1 + p_2$
3	0.365	$p_0 + p_1 + p_2 + p_3$
4	0.487	$p_0 + p_1 + p_2 + p_3 + p_4$

Fig. 3 plots the results for both experiments. Notice that, in both experiments, thresholds are highly heterogeneous and exhibit significant levels of non-compliance. In addition, since the inferred distributions of thresholds necessarily generate the observed results, they can be used to rationalise any differences observed in treatment effects across the two settings. For example, the observed non-linearity in the mask experiment can be rationalised by postulating that an especially large fraction have a threshold of 4.




Fig. 3. Threshold distributions. Notes. This figure shows the distributions of individual thresholds calculated from our two sets of experimental estimates. Individuals with thresholds of 0 and 5 are ‘always doers’ and ‘never doers’ respectively.

Estimating the long-run. Having discussed what might generate the estimates that we observe, we turn to the question about what the estimates might suggest about long-run behaviour. Answering this question requires making additional assumptions that we formalise using a dynamic model.²⁰ Since these assumptions are contestable (as discussed below), one should not place too much weight on the exact quantitative predictions that the model generates. Nonetheless, the model can cast some light on the more qualitative question of whether behaviour will converge to a single norm in the long run (a key question studied by, e.g., Young, 1993; Centola and Baronchelli, 2015, and others).

To answer this question, we study an evolutionary model that is described in Appendix B. In the model, agents must choose whether to take an action and can observe the choices of each of their four neighbours (note that both of these assumptions match up with our experimental environments). Agents make this decision by comparing their threshold with the number of their

¹⁹ This approach implicitly treats the set of possible thresholds as discrete. It is also possible, however, to view (normalised) thresholds as continuously distributed on [0, 1]. In that case, our estimates can be used to calculate the share of normalised thresholds that are zero, the share that are between 0 and 0.25, etc.

²⁰ This model is in the spirit of Granovetter (1978) and Schelling (1971), who also consider agents who repeatedly choose based on the share of others who are already taking the action. Unlike Granovetter (1978), our agents react only to those in their immediate environment (as in our experiments). Unlike Schelling (1971), our agents repeatedly revise a binary action (instead of repeatedly changing their location within the network).

neighbours already taking the action. To calibrate the model, we assume that the distribution of thresholds is that obtained from our experimental estimates.

Tables B.18 and B.19 reveal the model's predictions. Once calibrated using the threshold distribution from the masks experiment, the model predicts that around 23%–24% will wear a mask in the long-run (with almost no dependence on the initial conditions). Once calibrated using the results from the cameras experiment, the model predicts that 31%–36% will ultimately choose to use their camera. In these equilibria, a fraction of the adopters choose to adopt because they do this regardless of their neighbours' decisions; the remainder do so due to copying-like behaviour.

There are several reasons why the exact numbers generated by our model should not be taken too literally. First, the model assumes that the distribution of thresholds is fixed over time; in reality, this need not be the case. Second, the model's calibration assumes that agents react to their neighbours in the way that they react in our experimental settings. In reality, however, subjects might have viewed the experimenters as 'authority figures' and thus reacted more strongly to their behaviour than they would in more common settings.²¹ Third, the model assumes that individuals interact in a very stylised way (e.g. since all agents react to four neighbours). While it is possible to relax this by allowing agents to react to different numbers of neighbours, this requires assuming that our estimates are 'scale invariant' (e.g. that 1/4 individuals doing the action is the 'same' as 2/8) and using interpolation to estimate the shares of the 'missing' thresholds.

Although the quantitative predictions are thus not exact, the general qualitative prediction is more robust: despite the existence of the copying-like behaviour, the system does not converge to a situation in which either everybody or nobody does the behaviour. Intuitively, this is because our experimental estimates point to substantial heterogeneity in how individuals respond to social pressure. If instead individuals responded identically (using a common threshold), then the system would indeed 'snowball' to a single norm as predicted by canonical models of evolutionary game theory (see, e.g., Young 1993).

5. Concluding remarks

In this paper, we conduct multi-treatment experiments to obtain a quantitative understanding of how individuals' behaviour varies with the share doing an action in their immediate environment. Despite some differences between the estimates across our contexts, we obtain many commonalities across the two experiments. Perhaps most importantly, both sets of results point to substantial non-compliance, which emphasises the limits (along with the power) of social pressure in our settings.

Despite the large number of social norm experiments, we believe that our findings open up several avenues for future research. First, it may be worthwhile to conduct more experiments with multi-treatment randomisation in additional contexts. In particular, this could provide further evidence on whether the non-compliance that we document is robust. Second, it may be worthwhile to conduct such experiments with an even larger number of treatment groups, thus allowing for an even richer understanding of how individuals respond to social pressure. Given the very large sample sizes required to do this, however, such experiments are likely to be even more logistically challenging to implement than the two field experiments whose results we report here.

Declaration of competing interest

We confirm that we have no conflicts of interest to declare.

Data availability

We have included a link to our replication package on the title page.

²¹ Having said this, we should stress that the research assistants involved in the masks study were fellow undergraduates, and that the study took place in a relatively informal environment (namely, a room in their college). Moreover, while genuine peer effects might be stronger (since subjects dislike disappointing their friends), they might also be weaker (since friends are not authority figures).

Appendix A. Tables and figures

See Tables A.1–A.17.

Table A.1
Sample allocation (experiment 1).

Treatment	Frequency	Percentage
0	127	19.7
1	134	20.7
2	128	19.8
3	124	19.2
4	133	20.6
Total	646	100.0

Notes. This table shows how many subjects were allocated into each of the five treatments in the first experiment.

Table A.2
Descriptive statistics (experiment 1).

Variable	Mean	Std. Dev.
Age	20.8	3.90
Male	.497	.500
Humanities	.283	.451
MPLS	.240	.427
Medical sciences	.127	.333
Social sciences	.333	.471
Pre	.201	.401
<i>n</i>	646	

Notes. This table shows the descriptive statistics for experiment 1 (see Table 1 for a description of the variables).

Table A.3
Logit regressions (experiment 1).

Variable	No controls	Main specification	All controls
Treatment 1	.044 [.047]	.033 [.030]	.029 [.034]
Treatment 2	.171*** [.053]	.073** [.032]	.079** [.035]
Treatment 3	.238*** [.055]	.162*** [.040]	.168*** [.043]
Treatment 4	.331*** [.054]	.283*** [.042]	.304*** [.046]
Pre		.504*** [.030]	.498*** [.031]
Age		.003 [.005]	.002 [.004]
Male		-.006 [.026]	-.002 [.028]
<i>n</i>	646	646	620

Notes. This table reports the exact same specifications reported on in Table 2, except using logistic instead of linear regressions. Robust standard errors in parentheses (** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table A.4
Probit regressions (experiment 1).

Variable	No controls	Main specification	All controls
Treatment 1	.044 [.047]	.036 [.031]	.029 [.034]
Treatment 2	.171*** [.053]	.078** [.033]	.078** [.035]
Treatment 3	.238*** [.055]	.163*** [.040]	.162*** [.043]
Treatment 4	.331*** [.054]	.284*** [.043]	.298*** [.046]
Pre		.518*** [.024]	.512*** [.027]
Age		.002 [.004]	.001 [.004]
Male		-.007 [.026]	-.004 [.028]
<i>n</i>	646	646	620

Notes. This table reports the exact same specifications reported on in Table 2, except using probit instead of linear regressions. Robust standard errors in parentheses (** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table A.5
Comparisons (experiment 1).

Comparison	No controls	Main specification	All controls
T0 vs. T1	.355	.278	.545
T1 vs. T2	.019	.205	.160
T2 vs. T3	.269	.051	.069
T3 vs. T4	.131	.019	.018
T0 vs. T2	.001	.014	.032
T1 vs. T3	.001	.002	.002
T2 vs. T4	.008	.000	.000
T0 vs. T3	.000	.000	.000
T1 vs. T4	.000	.000	.000
T0 vs. T4	.000	.000	.000

Notes. This table reports p -values corresponding to hypothesis that the effect of treatment k is the same as the effect of treatment k' , for all possible $k \neq k'$. We do this for the three specifications considered in Table 2.

Table A.6
Regressions conditional on pre (experiment 1).

Variable	Pre = 0	Pre = 1
Treatment 1	.050* [.030]	-.091 [.100]
Treatment 2	.079** [.036]	-.005 [.067]
Treatment 3	.197*** [.046]	-.014 [.071]
Treatment 4	.339*** [.050]	.018 [.067]
Age	.003 [.006]	-.001 [.003]
Male	-.015 [.031]	.043 [.045]
Constant	-.023 [.134]	.946*** [.072]
n	516	130
R^2	.115	.025

Notes. This table shows the regressions from the main specification (see Table 2) conditional on whether subjects entered the room with or without a mask (corresponding to Pre = 1 and Pre = 0 respectively). Robust standard errors in parentheses (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.7
Changes (experiment 1).

	T0	T1	T2	T3	T4
Putting mask on	.028	.080	.106	.223	.368
Taking mask off	.056	.143	.059	.067	.037

Notes. The first row shows the share who put a mask on given that they entered the room without wearing a mask. The second row shows the share who took their mask off given that they entered the room wearing a mask.

Table A.8
Polynomial regressions (experiment 1).

Variable	Linear	Quadratic	Cubic
Masks	.070*** [−.010]	0.008 [−.028]	0.024 [−.062]
Masks ²		.016** [−.008]	.004 [−.045]
Masks ³			.002 [−.008]
Pre	.752*** [−.029]	.757*** [−.029]	.757*** [−.029]
Age	.002 [−.005]	.002 [−.005]	.002 [−.005]
Male	−.008 [−.026]	−.007 [−.026]	−.007 [−.026]
Constant	−.022 [−.102]	.016 [−.107]	.014 [−.107]
Joint test	.000	.000	.000
R ²	.491	.494	.494

Notes. In this table, we regress whether subjects chose to wear a mask on the number of experimenters wearing a mask, as well higher order terms to capture potential non-linearity (we also control for ‘pre’, age, and gender). The penultimate row reports p -values corresponding to the hypothesis that the coefficients on all mask variables are zero.

Table A.9
Explanations from online survey.

Explanation	Frequency
Trying to avoid judgement	.148
Trying to cater to others’ preferences	.511
Trying to follow rules	.148
Reciprocity	.023
COVID risks	.011
Not answering question	.159
n	88

Notes. This table shows the frequencies of the explanations given by subjects (see Appendix D for a detailed description of the categories).

Table A.10
Sample allocation (experiment 2).

Treatment	Frequency	Percentage
0	232	20.8
1	204	18.3
2	223	20.0
3	241	21.7
4	213	19.1
Total	1113	100.0

Notes. This table shows how many subjects were allocated into each of the five treatments in the second experiment.

Table A.11
Descriptive statistics (experiment 2).

Variable	Mean	Std. Dev.
Age	42.4	13.9
Male	.465	.499
n	1113	

Notes. This table shows the descriptive statistics for experiment 2.

Table A.12
Logit regressions (experiment 2).

Variable	No controls	Main specification	All controls
Treatment 1	.077* [.043]	.127*** [.039]	.133*** [.039]
Treatment 2	.176*** [.043]	.215*** [.039]	.218*** [.040]
Treatment 3	.281*** [.043]	.314*** [.039]	.323*** [.045]
Treatment 4	.355*** [.044]	.385*** [.041]	.389*** [.051]
Pre		.741*** [.092]	.743*** [.092]
Age		.000 [.001]	.000 [.001]
Male		.023 [.027]	.023 [.027]
<i>n</i>	1113	1111	1109

Notes. This table reports the exact same specifications reported on in Table 4, except using logistic instead of linear regressions. Robust standard errors in parentheses (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table A.13
Probit regressions (experiment 2).

Variable	No controls	Main specification	All controls
Treatment 1	.077* [.043]	.125*** [.039]	.130*** [.039]
Treatment 2	.176*** [.043]	.216*** [.039]	.218*** [.040]
Treatment 3	.281*** [.043]	.312*** [.039]	.321*** [.046]
Treatment 4	.355*** [.044]	.385*** [.040]	.389*** [.052]
Pre		.701*** [.075]	.699*** [.076]
Age		.000 [.001]	.000 [.001]
Male		.024 [.027]	.025 [.027]
<i>n</i>	1113	1111	1109

Notes. This table reports the exact same specifications reported on in Table 4, except using probit instead of linear regressions. Robust standard errors in parentheses (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table A.14
Comparisons (experiment 2).

Comparison	No controls	Main specification	All controls
T0 vs. T1	.074	.003	.002
T1 vs. T2	.035	.043	.051
T2 vs. T3	.022	.028	.020
T3 vs. T4	.116	.116	.152
T0 vs. T2	.000	.000	.000
T1 vs. T3	.000	.000	.000
T2 vs. T4	.000	.000	.001
T0 vs. T3	.000	.000	.001
T1 vs. T4	.000	.000	.001
T0 vs. T4	.000	.000	.001

Notes. This table reports p -values corresponding to hypothesis that the effect of treatment k is the same as the effect of treatment k' , for all possible $k \neq k'$. We do this for the three specifications considered in Table 4.

Table A.15
Regressions conditional on pre (experiment 2).

Variables	Pre = 0	Pre = 1
Treatment 1	0.136*** [0.042]	-0.007 [0.129]
Treatment 2	0.225*** [0.042]	0.113* [0.060]
Treatment 3	0.334*** [0.042]	0.038 [0.080]
Treatment 4	0.407*** [0.044]	0.110* [0.063]
Age	0.000 [0.001]	-0.002 [0.002]
Male	0.025 [0.029]	-0.026 [0.050]
Constant	0.136*** [0.049]	0.977*** [0.120]
<i>n</i>	1031	80
<i>R</i> ²	0.088	0.053

Notes. This table shows the regressions from the main specification (see Table 4) conditional on whether subjects entered the room with or without their camera on (corresponding to Pre = 1 and Pre = 0 respectively). Robust standard errors in parentheses (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table A.16
Changes (experiment 2).

	T0	T1	T2	T3	T4
Turning camera on	0.156	0.296	0.381	0.491	0.566
Turning camera off	0.111	0.125	0.000	0.059	0.000

Notes. The first row shows the share who turned their camera on given that they joined the call without video. The second row shows the share who turned their camera off given that they joined the call with video.

Table A.17
Polynomial regressions (experiment 2).

Variable	Linear	Quadratic	Cubic
Cameras	.095*** [-.009]	.119*** [-.032]	.119 [-.074]
Cameras ²		-.006 [-.008]	-.006 [-.049]
Cameras ³			.000 [-.008]
Pre	.576*** [-.033]	.578*** [-.033]	.578*** [-.033]
Age	.000 [-.001]	.000 [-.001]	.000 [-.001]
Male	.023 [-.027]	.024 [-.027]	.024 [-.027]
Constant	.169*** [-.046]	.156*** [-.047]	.156*** [-.047]
Joint test	.000	.000	.000
<i>R</i> ²	.161	.161	.161

Notes. In this table, we regress whether subjects chose to use their camera on the number of experimenters using a camera, as well higher order terms to capture potential non-linearity (we also control for 'pre', age, and gender). The penultimate row reports p -values corresponding to the hypothesis that the coefficients on all camera variables are zero.

Appendix B. An evolutionary model

In this section, we outline an evolutionary model of local interaction in overlapping networks. The model presented here shares some similarities to the model studied by Efferson et al. (2020). An important difference is that, while Efferson et al. (2020) assume that decision makers choose randomly, we instead assume that they choose deterministically but with heterogeneous decision rules. In addition, our model assumes that individuals respond to the decisions of their ‘neighbours’ (in line with our experimental settings); whereas Efferson et al. (2020) assume that they best respond to the entire population.

In our baseline model, we assume the following:²²

- There are l^2 agents, each located on a node of a grid with side length $l \in \mathbb{N}^+$. Let (r, c) denote the agent located at row r and column c ; so the set of agents is the set $N = \{(r, c) : r \in \{1, \dots, l\}, c \in \{1, \dots, l\}\}$.
- As in our experiments, agents are faced with a binary choice: they must either take an action (denoted $a_{r,c} = 1$) or not take the action (denoted $a_{r,c} = 0$).
- Each agent (r, c) has a set of ‘neighbours’ $N_{r,c}$ whose actions they can see. For each agent, we assume that $N_{r,c} = \{(i, j) : (i, j) \in N, |i - r| \leq 1, |j - c| \leq 1, (i, j) \neq (r, c)\}$. Observe that agents in the interior have 4 neighbours, agents on the edge have 3 neighbours, and agents in the corners have 2 neighbours.
- We define $m_{r,c}^1$ as the share of individual (r, c) ’s neighbours who have chosen to do the action. Formally, $m_{r,c}^1 = \frac{1}{|N_{r,c}|} \sum_{(i,j) \in N_{r,c}} a_{i,j}$ where $|N_{r,c}|$ is the cardinality of $N_{r,c}$.
- Each agent is endowed with a (fixed) threshold $\tau_{r,c} \in [0, 1]$. As in the main text, we assume that they choose $a_{r,c} = 1$ if and only if $m_{r,c}^1 \geq \tau_{r,c}$.
- Agents interact over multiple periods. In each period, one agent is chosen to move at random; and updates their action (if necessary) by comparing their threshold $\tau_{r,c}$ with the share of their neighbours who are taking the action $m_{r,c}^1$.

To assess the robustness of our results, we also study an alternative model that departs from the model sketched above in various ways.²³ In this model — which we label the *edgeless model* — each agent is linked with the same number of neighbours. In addition, each agent has a probability $\epsilon \in [0, 1]$ of making a ‘mistake’, i.e. choosing the opposite action as that required by their threshold. Finally, a share $p \in [0, 1]$ of agents are selected in period to revise their action; so in principle multiple agents can update their action simultaneously.

To simulate the results of our models, we use the following procedure:

- We specify a distribution of thresholds in the population, and randomly scatter these thresholds across the agents.
- We also specify the share of agents who initially take the action; and we randomly scatter the agents who are taking the action on grid.
- We then allow the model to run for 1000 periods (or until it is ‘stable’ so no further changes can occur).

As discussed above, the results of the model could in principle depend on the way in which thresholds and initial actions are scattered. As a result, we conduct all simulations 1000 times and report the distribution of results across simulations.

Before turning to our main results, we provide a simple example to illustrate the mechanics of the model. To generate this example, we suppose that, initially, 40% of agents are taking the action; and we set $s = 5$. In addition, we assume (for expositional simplicity) that all agents have a threshold $\tau_{r,c} = 0.5$, so choose $a_{r,c} = 1$ if and only if half or more of their neighbours are doing the action. While one would normally repeat the simulation many times, here we just report the outcome of one simulation.

After randomly scattering the initial actions, we obtain the initial state

0	1	1	1	0
0	1	0	0	1
1	0	1	1	0
0	0	0	0	0
1	0	0	1	0

As can be seen, 10 of the $s^2 = 25$ agents initially take the action (indicated by a 1); the rest do not. Several rounds now progress in which the player chosen to move does not wish to update their action. Eventually, however, the player at row 3 and column 1 is chosen to move (they are coloured in red). Since none of their neighbours (coloured in blue) were taking the action, they choose to switch to action 0. This yields the new state

²² The corresponding code can be viewed here: https://github.com/Itzhak95/tipping_points.

²³ The corresponding code can be viewed here: https://github.com/rrozzi/tipping_point-netlogo.

0	1	1	1	0
0	1	0	0	1
0	0	1	1	0
0	0	0	0	0
1	0	0	1	0

As the process continues, additional players are given the opportunity to also revise their action. After 13 such revisions, we finally obtain the state

1	1	0	0	0
1	1	0	0	0
0	0	0	0	0
0	0	0	0	0
0	0	0	0	0

This state is *stable* in the sense that no agent has an incentive to change their behaviour. The agent in the top left is surrounded by neighbours who choose $a_{r,c} = 1$, so would also want to choose $a_{r,c} = 1$ if allowed to update their action. The agent at (2, 2) is surrounded by 4 neighbours, half of whom are taking the action; so also chooses $a_{r,c} = 1$ (recall that all thresholds are set at $\tau_{r,c} = 0.5$). Meanwhile, the agents at (2, 1) and (1, 2) are each surrounded by 3 neighbours, 2 of whom are choosing the action; so they also wish to choose the action. Finally, one can verify that the agents choosing $a_{r,c} = 0$ are choosing optimally given their threshold and the share of their neighbours who are taking the action.

We now calibrate our model using the threshold distributions calculated in Section 4. We assume a population size of 100; and the edgeless model further assumes an error probability $\epsilon = 0.01$ and a probability of revision $p = 0.07$. As stated above, each simulation is run for 1000 periods (or until the obtained state is stable); and all simulations are conducted 1000 times. Tables B.18 and B.19 display the results for experiment 1 (face masks) and experiment 2 (Zoom calls) respectively. The first row specifies the initial share who are assumed to do the action. The rows ‘mean (main)’ and ‘mean (edgeless)’ display the average share who end up doing the activity in the main specification and edgeless model respectively. The rows ‘Var (main)’ and ‘Var (edgeless)’ specify the variance of outcomes across simulations.

Table B.18
Simulation results (experiment 1).

Initial share	0.05	0.15	0.25	0.35	0.45	0.55	0.65	0.75	0.85	0.95
Mean (main)	.228	.230	.230	.231	.231	.233	.234	.235	.235	.236
Var (main)	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000
Mean (edgeless)	.237	.242	.242	.242	.240	.242	.240	.239	.240	.239
Var (edgeless)	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000

Notes. This table shows the results of simulating our models using the distribution of thresholds obtained by experiment 1 (see Table 2). That is, we set $p_0 = .203$, $p_1 = .033$, $p_2 = .044$, $p_3 = .085$, $p_4 = .123$, $p_5 = .513$.

Table B.19
Simulation results (experiment 2).

Initial share	0.05	0.15	0.25	0.35	0.45	0.55	0.65	0.75	0.85	0.95
Mean (main)	.311	.316	.321	.329	.333	.338	.341	.345	.349	.353
Var (main)	.001	.001	.001	.001	.001	.001	.001	.001	.001	.001
Mean (edgeless)	.353	.353	.347	.356	.355	.362	.354	.360	.356	.355
Var (edgeless)	.001	.001	.001	.001	.001	.001	.001	.001	.001	.001

Notes. This table shows the results of simulating our models using the distribution of thresholds obtained by experiment 2 (see Table 4). That is, we set $p_0 = .209$, $p_1 = .118$, $p_2 = .091$, $p_3 = .099$, $p_4 = .072$, $p_5 = .411$.

Three results are apparent. First, we see that the results of the simulations are relatively insensitive to the initial share who are assumed to do the activity. This is especially true in the edgeless model since this assumes that agents occasionally make errors, which weakens dependence on initial conditions in the usual way (Young, 1993). Second, the variance in outcomes across simulations is

very low, which is again points to the lack of importance of initial conditions (since different simulations generate different outcomes only due to variation in initial conditions). Finally, and most importantly, we see that our models generate convergence to interior equilibria. This should not come as a surprise given that some agents always do the action, that other agents never do the action, and that a final group of agents engage in copying behaviour.

Appendix C. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jebo.2024.106699>.

References

- Agerström, J., Carlsson, R., Nicklasson, L., Guntell, L., 2016. Using descriptive social norms to increase charitable giving: The power of local norms. *J. Econ. Psychol.* 52, 147–153.
- Allcott, H., 2011. Social norms and energy conservation. *J. Public Econ.* 95 (9–10), 1082–1095.
- Allcott, H., Kessler, J.B., 2019. The welfare effects of nudges: A case study of energy use social comparisons. *Am. Econ. J.: Appl. Econ.* 11 (1), 236–276.
- Andreoni, J., Nikiforakis, N., Siegenthaler, S., 2021. Predicting social tipping and norm change in controlled experiments. *Proc. Natl. Acad. Sci.* 118 (16), e2014893118.
- Angrist, J.D., Imbens, G.W., Rubin, D.B., 1996. Identification of causal effects using instrumental variables. *J. Amer. Statist. Assoc.* 91 (434), 444–455.
- Asch, S.E., 1951. Effects of group pressure upon the modification and distortion of judgments. *Organ. Influ. Process.* 58, 295–303.
- Ayres, I., Raseman, S., Shih, A., 2013. Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage. *J. Law Econ. Organ.* 29 (5), 992–1022.
- Barceló, J., Sheen, G.C.-H., 2020. Voluntary adoption of social welfare-enhancing behavior: Mask-wearing in Spain during the COVID-19 outbreak. *PLOS ONE* 15 (12), e0242764.
- Barile, J.P., Guerin, R.J., Fisher, K.A., Tian, L.H., Okun, A.H., Vanden Esschert, K.L., Jeffers, A., Gurbaxani, B.M., Thompson, W.W., Prue, C.E., 2021. Theory-based behavioral predictors of self-reported use of face coverings in public settings during the COVID-19 pandemic in the United States. *Ann. Behav. Med.* 55 (1), 82–88.
- Bedenlier, S., Wunder, I., Gläser-Zikuda, M., Kammerl, R., Kopp, B., Ziegler, A., Händel, M., 2021. Generation invisible? Higher education students' (non) use of webcams in synchronous online learning. *Int. J. Educ. Res. Open* 2, 100068.
- Bicchieri, C., 2016. *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*. Oxford University Press.
- Bokemper, S.E., Cucciniello, M., Rotesi, T., Pin, P., Malik, A.A., Willebrand, K., Paintsil, E.E., Omer, S.B., Huber, G.A., Melegaro, A., 2021. Experimental evidence that changing beliefs about mask efficacy and social norms increase mask wearing for COVID-19 risk reduction: Results from the United States and Italy. *PLOS ONE* 16 (10), e0258282.
- Bond, R., Smith, P.B., 1996. Culture and conformity: A meta-analysis of studies using Asch's (1952b, 1956) line judgment task. *Psychol. Bull.* 119 (1), 111.
- Borsari, B., Carey, K.B., 2003. Descriptive and injunctive norms in college drinking: a meta-analytic integration. *J. Stud. Alcohol* 64 (3), 331–341.
- Bott, K.M., Cappelen, A.W., Sørensen, E.Ø., Tungodden, B., 2020. You've got mail: A randomized field experiment on tax evasion. *Manage. Sci.* 66 (7), 2801–2819.
- Bursztyl, L., Ederer, F., Ferman, B., Yuchtman, N., 2014. Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica* 82 (4), 1273–1301.
- Cason, T.N., Mui, V.-L., 1998. Social influence in the sequential dictator game. *J. Math. Psych.* 42 (2–3), 248–265.
- Castelli, F.R., Sarvary, M.A., 2021. Why students do not turn on their video cameras during online classes and an equitable and inclusive plan to encourage them to do so. *Ecol. Evol.* 11 (8), 3565–3576.
- Centola, D., Baronchelli, A., 2015. The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proc. Natl. Acad. Sci.* 112 (7), 1989–1994.
- Centola, D., Becker, J., Brackbill, D., Baronchelli, A., 2018. Experimental evidence for tipping points in social convention. *Science* 360 (6393), 1116–1119.
- Cialdini, R.B., 2007. Descriptive social norms as underappreciated sources of social control. *Psychometrika* 72 (2), 263–268.
- Cialdini, R.B., Reno, R.R., Kallgren, C.A., 1990. A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *J. Pers. Soc. Psychol.* 58 (6), 1015.
- Costa, D.L., Kahn, M.E., 2013. Energy conservation “nudges” and environmentalist ideology: Evidence from a randomized residential electricity field experiment. *J. Eur. Econom. Assoc.* 11 (3), 680–702.
- Damm, A.P., Dustmann, C., 2014. Does growing up in a high crime neighborhood affect youth criminal behavior? *Amer. Econ. Rev.* 104 (6), 1806–1832.
- Efferson, C., Vogt, S., Elhadi, A., Ahmed, H.E.F., Fehr, E., 2015. Female genital cutting is not a social coordination norm. *Science* 349 (6255), 1446–1447.
- Efferson, C., Vogt, S., Fehr, E., 2020. The promise and the peril of using social influence to reverse harmful traditions. *Nat. Hum. Behav.* 4 (1), 55–68.
- Ehret, S., Constantino, S., Weber, E., Efferson, C., Vogt, S., 2022. Group identities make fragile tipping points.
- Farrow, K., Grolleau, G., Ibanez, L., 2017. Social norms and pro-environmental behavior: A review of the evidence. *Ecol. Econom.* 140, 1–13.
- Ferraro, P.J., Price, M.K., 2013. Using nonpecuniary strategies to influence behavior: evidence from a large-scale field experiment. *Rev. Econ. Stat.* 95 (1), 64–73.
- Fortin, B., Lacroix, G., Villeval, M.-C., 2007. Tax evasion and social interactions. *J. Public Econ.* 91 (11–12), 2089–2112.
- Freidin, E., Acera Martini, L., Senci, C.M., Duarte, C., Carballo, F., 2022. Field observations and survey evidence to assess predictors of mask wearing across different outdoor activities in an Argentine city during the COVID-19 pandemic. *Appl. Psychol.: Health Well-Being* 14 (1), 81–100.
- Friend, R., Rafferty, Y., Bramel, D., 1990. A puzzling misinterpretation of the Asch ‘conformity’ study. *Eur. J. Soc. Psychol.* 20 (1), 29–44.
- Gerber, A.S., Rogers, T., 2009. Descriptive social norms and motivation to vote: Everybody's voting and so should you. *J. Politics* 71 (1), 178–191.
- Gherheș, V., Șimon, S., Para, I., 2021. Analysing students' reasons for keeping their webcams on or off during online classes. *Sustainability* 13 (6), 3203.
- Goldberg, M., Gustafson, A., Maibach, E., van der Linden, S., Ballew, M.T., Bergquist, P., Kotcher, J., Marlon, J.R., Rosenthal, S., Leiserowitz, A., 2020. Social norms motivate COVID-19 preventive behaviors. *PsyArXiv*.
- Goldstein, N.J., Cialdini, R.B., Griskevicius, V., 2008. A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *J. Consum. Res.* 35 (3), 472–482.
- Granovetter, M., 1978. Threshold models of collective behavior. *Am. J. Sociol.* 83 (6), 1420–1443.
- Heldt, T., 2005. Conditional cooperation in the field: cross-country skiers' behavior in Sweden. In: *Economic Science Association, European Meeting, Alessandria, Italian*, September 15–18, 2005.
- Henrich, J., Heine, S.J., Norenzayan, A., 2010. The weirdest people in the world? *Behav. Brain Sci.* 33 (2–3), 61–83.
- Ichino, A., Maggi, G., 2000. Work environment and individual background: Explaining regional shirking differentials in a large Italian firm. *Q. J. Econ.* 115 (3), 1057–1090.
- Jackson, M.O., Yariv, L., 2007. Diffusion of behavior and equilibrium properties in network games. *Amer. Econ. Rev.* 97 (2), 92–98.
- Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation, and long run equilibria in games. *Econometrica* 29–56.
- Kessler, J.B., Low, C., Singhal, M., 2021. Social policy instruments and the compliance environment. *J. Econ. Behav. Organ.* 192, 248–267.

- Kreindler, G.E., Young, H.P., 2013. Fast convergence in evolutionary equilibrium selection. *Games Econom. Behav.* 80, 39–67.
- Krupka, E., Weber, R.A., 2009. The focusing and informational effects of norms on pro-social behavior. *J. Econ. Psychol.* 30 (3), 307–320.
- Lefebvre, M., Pestieau, P., Riedl, A., Villeval, M.C., 2015. Tax evasion and social information: an experiment in Belgium, France, and the Netherlands. *Int. Tax Public Finance* 22 (3), 401–425.
- Linek, M., Traxler, C., 2021. Framing and social information nudges at Wikipedia. *J. Econ. Behav. Organ.* 188, 1269–1279.
- Manski, C.F., 1993. Identification of endogenous social effects: The reflection problem. *Rev. Econ. Stud.* 60 (3), 531–542.
- Martin, R., Randal, J., 2008. How is donation behaviour affected by the donations of others? *J. Econ. Behav. Organ.* 67 (1), 228–238.
- Mascagni, G., 2018. From the lab to the field: A review of tax experiments. *J. Econ. Surv.* 32 (2), 273–301.
- Novak, L., 2020. Persistent norms and tipping points: The case of female genital cutting. *J. Econ. Behav. Organ.* 177, 433–474.
- Rudert, S.C., Janke, S., 2021. Following the crowd in times of crisis: Descriptive norms predict physical distancing, stockpiling, and prosocial behavior during the COVID-19 pandemic. *Group Process. Intergroup Relat.* 13684302211023562.
- Schelling, T.C., 1971. Dynamic models of segregation. *J. Math. Sociol.* 1 (2), 143–186.
- Sederevičiūtė-Pačiauskienė, Ž., Valantinaitė, I., Asakavičiūtė, V., 2022. ‘Should I turn on my video camera?’ the students’ perceptions of the use of video cameras in synchronous distant learning. *Electronics* 11 (5), 813.
- Smith, S., Windmeijer, F., Wright, E., 2015. Peer effects in charitable giving: Evidence from the (running) field. *Econ. J.* 125 (585), 1053–1071.
- Thöni, C., Gächter, S., 2015. Peer effects and social preferences in voluntary cooperation: A theoretical and experimental analysis. *J. Econ. Psychol.* 48, 72–88.
- Woodcock, A., Schultz, P.W., 2021. The role of conformity in mask-wearing during COVID-19. *PLOS ONE* 16 (12), e0261321.
- Young, H.P., 1993. The evolution of conventions. *Econometrica* 57–84.
- Young, H.P., 2009. Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *Amer. Econ. Rev.* 99 (5), 1899–1924.