Identifying Social Norms: Theory and Evidence from Prosocial Behaviour

Toke Aidt Julia Shvets

January 15, 2026

Abstract

When someone decides whether to take a prosocial action, they might be influenced by what others around them do. This could happen for two reasons. The first is standard externalities: for instance, if one person's prosocial decisions lowers the cost of doing the same for another person. The second are social norms – people may want to conform to the decisions of others. A well known non-identifiability result shows that standard peer effect estimations cannot separate these two models (Brock & Durlauf 2001, Blume et al 2015, Boucher & Fortin 2016), despite their fundamentally different motives and often divergent policy implications (Young 2015, Bhattacharya et al 2024). We develop a method to overcome this problem and identify social norms in observational data. Using a theoretical model of choices in groups, we show that potential multiplicity of social norms can create incentives for behaviour that are very different from those under externalities. These differences, although not identifiable in peer effect regressions, can be identified by looking at the full distribution of choices across groups. We develop this method and identify social norms in the context of students' voluntary participation in Cambridge's large-scale asymptomatic Covid-19 testing programme. We show that students develop opposing social norms in different households. This heterogeneity is concealed by both peer effect estimates and aggregate testing patterns, whilst it can reverse some of the policy conclusions that flow from them.

1 Introduction

Many socially important actions – voting, helping others, donating to charity – require individuals to incur a private cost in order to deliver a social benefit. Often, people engage in such prosocial actions when others in their peer group do the same. This can happen for two reasons. On the one hand, there can be standard externalities. For example, if one friend has found a volunteering opportunity, this lowers the cost for another friend to join in (they do not need to do the research) and thus makes it more likely that they will also volunteer. On the other hand, an individual may want to volunteer when others around him do so because conforming to the choices of his peers makes him feel better. In other words, the decision of whether to take the prosocial action depends on the prevailing social norm in the peer group.[1]

Both of these mechanisms – externalities and norms – imply that choices within a peer group are strategic complements. As a consequence, standard econometric models of peer effects cannot separate norms from externalities (Brock and Durlauf 2001; Blume et al. 2015; Boucher et al. 2014; Boucher et al. 2024). In this paper, we develop a novel approach to distinguish social norms from externalities in observational data on prosocial actions. Doing so is important, as the two mechanisms correspond to fundamentally different models of human behaviour (see, e.g., Elster 1989; Bernheim 1994; Bicchieri et al. 2023) and carry different policy implications (see, e.g., Young 2015; Bhattacharya, Dupas, and Kanaya 2024).

We analyse peer influences and separate the effects of social norms from those of externalities in the context of young people's behaviour during the Covid pandemic. At its onset, the University of Cambridge set up a large-scale asymptomatic screening programme that asked 16,000 of its students[2] to voluntarily take a weekly Covid-19 test. The goal of the programme, as emphasised in its recruitment material, was to 'keep Cambridge safe' by detecting infections among asymptomatic carriers and thereby limiting their spread (University of Cambridge 2020). Alongside these social benefits, taking a test without symptoms also involves non-negligible private costs: it is unpleasant, inconvenient, and carries the risk of self-isolation or quarantine. Therefore, the decision to test is a clear example of a prosocial action (Levitt 2020; Fang et al. 2022).

During the pandemic period, Cambridge students were assigned to 'households' groups, with six students on average, consisting of student rooms that shared a kitchen, common

1. For a survey of the substantial literature on social norms, see Bicchieri et al. (2023).
2. This includes all students living in college accommodation, roughly two thirds of the total student body.

room or other facilities. This household structure creates an environment in which both household-level peer effects would naturally emerge and in which both social norms and payoff externalities may influence the students' testing decisions.

First, observability of testing decisions and frequent social interactions inside a household create the potential for household-level social norms to develop. Moreover, different norms can plausibly emerge. On the one hand, some students may believe that testing is 'the right thing to do' and expect this of the others, developing a pro-testing norm. On the other hand, the rules of the programme require all members of a household to quarantine if any one member tests positive. As a result, some households may develop an anti-testing norm, where asymptomatic testing would be met with disapproval since it could bring quarantine upon everyone in the household. In sum, some households may develop a norm in which members are expected to test, whilst the opposite norm, that of not testing, may emerge in others. Second, the programme also generates payoff externalities within households. Because all tests for a household must be picked up and dropped off together in a single envelope,[3] a single student can undertake this task on behalf of the entire household, thereby reducing the cost of testing for other members.

We use detailed data from the Cambridge asymptomatic screening programme on individual student's testing decisions to study peer effects, and we develop a novel test that enables us to demonstrate that social norms play an important role. To make the logic underlying the test clear and to explain how and when it is possible to separate norms from externalities in an environment with strong peer effects, we develop a theoretical model in which students assigned to household groups (their peer group) decide to test or not (i.e., to take the prosocial action or not). We consider three scenarios: (1) a student's payoff of testing is independent of the actions of others in the household (independent payoffs), (2) household members exert positive payoff externalities on each other (positive externalities), and (3) students get a benefit from conforming to a social norm as embodied by the behaviour of others in their household (social norms).

First, we show that both scenarios (2) and (3) result in strategic complementarity between students' actions and predict positive peer effects. Consequently, standard peer effect regressions cannot separate norms from externalities, a well-known result in the peer effects literature (Blume et al. 2015; Boucher and Fortin 2016).

Second, we show that scenarios (2) and (3) lead to qualitatively different predictions for

3. This is because tests are pooled to save on costs; see Section 2 for details.

how the distribution of the number of students who test across households differs from the benchmark case of independent testing decisions in scenario (1). Positive externalities create incentives for students who would not test under independence to do so, and thus shift the distribution towards more students testing. Importantly, since no student who tested in the absence of a positive externality would *stop* testing in its presence, this implies that the share of households in which all students decide not to test cannot be larger in the presence of a positive externality than in its absence. In contrast, the presence of a social norm to conform to what others in a household do can increase the share of households in which no one tests relative to scenario (1) with independent payoffs. Intuitively, if the household's social norm is not to test, it can induce students who would have tested in its absence, in order to conform to the behaviour of other housemates, to stop testing. It is this difference in how externalities and social norms affect the distribution of outcomes across households relative to independence that enables us to distinguish between them and to establish whether social norms are present.[4,5]

Our three main findings are as follows. First, consistent with both the presence of household-level externalities and social norms, we identify positive peer effects in students' testing decisions. The richness of our data enables us to rule out many potential confounding factors discussed in the literature (see, e.g., Manski 1993), using fixed effects and random assignment to households for a large fraction of students. The peer effects are large: when all other students in a household switch from not testing to testing, a student's probability of testing increases by around 70 percentage points. We also find positive peer effects on the extensive margin of enrolment in the screening program.

Second, we implement our test for norms versus externalities using the distributions of household-level outcomes. This requires a comparison of the observed distribution of the number of students testing across households with what it would have been if testing decisions were independent. To construct this counterfactual, we use the testing decisions of students who live alone in households of size one and therefore could not be affected by either externalities or social norms. We find that, compared to the constructed counterfactual, the actual distribution has significantly more mass at both extremes. In other words, instances in which everyone tests are more frequent, and instances in which no one tests are also more frequent, relative to what we would expect under independence. This implies that the observed student behaviour cannot be explained by externalities alone and establishes

4. As we show in Section 5.2, this is a sufficient test, i.e., social norms need to be strong enough for the test to detect them.

5. Negative externalities, which are unlikely in our context, are considered in Appendix C. We show that the test for social norms continues to work.

that social norms must have been present.[6]

Third, the data show that both types of social norms, namely norms of testing and norms of not testing, developed in different households.

Distinguishing between externalities and social norms matters not least because the two mechanisms have fundamentally different policy implications (Young 2015; Bhattacharya, Dupas, and Kanaya 2024). In the context of the asymptomatic testing programme, for example, we find that, on aggregate, testing is higher when students live with others than when they live alone. A natural interpretation of this pattern is that assigning students to groups is an effective way to increase testing. This conclusion is misleading if peer effects operate through social norms. The aggregate pattern conceals substantial heterogeneity across households, with some developing norms that encourage testing and others developing norms that discourage it. As a result, in the presence of social norms, grouping individuals can either increase or decrease aggregate testing, making the policy effect a priori ambiguous, as we show in our theoretical framework in Section 3.4. This contrasts sharply with a setting in which peer effects are driven by positive externalities, where grouping individuals unambiguously raises aggregate testing.

Our paper contributes to several strands of the literature.

First, it contributes to the literature on identifying social norms. This literature can be broadly divided into two methodological strands. The first relies on surveys and experiments that elicit beliefs about norms, such as the survey instruments developed by Bicchieri and Chavez (2010) or the incentive-compatible approach of Krupka and Weber (2013). Our approach is complementary to this strand in that it does not require direct access to subjects and is based on observed behaviour rather than stated beliefs.

The second strand exploits experimental (natural or randomized) variation, typically in observability or information, to generate behavioural changes that are difficult to explain other than by pressure to conform to a norm (see, e.g., the survey by Bursztyn and Jensen (2017)). While powerful, this approach may fail to detect social norms in environments with multiple equilibria, where similar groups can endogenously develop different norms (Bénabou and Tirole 2006; Postlewaite 2011; Young 2015)[7]. Our approach is complementary here as well: it is designed for observational data and is explicitly grounded in the idea that the

6. This finding is robust to the counterfactual being students who *ended up* alone in households because their housemates stayed with their parents during Covid. This addresses the possibility of self-selection.

7. For an exception, see (Bursztyn, Egorov, and Jensen 2019).

presence of multiple equilibria is a defining feature of environments governed by social norms.

Second, by developing a new approach to separating social norms from externalities, we contribute to the methodological literature on the identification of peer effects. This literature has long emphasised the observational equivalence of these two mechanisms in standard peer-effects regressions (e.g., Brock and Durlauf 2001; Blume et al. 2015; Boucher and Fortin 2016). Boucher and Fortin (2016) and Boucher et al. (2024) propose solutions based on modified peer-effects regression models. While both their approaches and ours require data on isolated individuals, there are important differences. First, existing approaches rely on substantial parametric and structural assumptions to achieve identification. Second, they impose a unique equilibrium on social interactions, in contrast to the view that multiple equilibria are a defining feature of social norms (e.g., Young 2015). By comparison, our approach imposes minimal structure on the data, and is based on differences in incentives for behaviour and, through them, distributions of outcomes, generated by social norms relative to externalities. Moreover, our model explicitly allows for multiple equilibria and exploits them for identification.

Third, this paper contributes to the broader literature on pro-social behaviour. Extensive evidence that individuals do not always act in accordance with narrow self-interest has led to the development of theories of social preferences (Fehr and Schmidt 1999). While these models have been highly influential and perform well in some settings (Gächter, Nosenzo, and Sefton 2013), they often fall short in explaining more complex behaviours in field environments (see DellaVigna (2009) for a survey). Our data show that pro-social behaviour is fundamentally shaped by group context: the likelihood of making a pro-social decision changes sharply when students are placed into households. This pattern cannot be accounted for by social preferences alone. Instead, we show that social norms play a central role. Our findings complement field evidence from natural and randomized experiments demonstrating the importance of social pressure in pro-social choices, including voting (Funk 2010; DellaVigna et al. 2017), charitable giving (DellaVigna, List, and Malmendier 2012; Smith, Windmeijer, and Wright 2015; Goette and Tripodi 2021), and work ethic (Mas and Moretti 2009). Crucially, unlike existing studies, our approach reveals that social interactions need not uniformly amplify pro-social behaviour. Different groups can develop opposing social norms, implying that social interactions can *reduce* pro-social behaviour by inducing individuals with pro-social preferences to conform to more selfish group norms.[8]

8. Related experimental work highlights the fragility of pro-social norms; see Bicchieri et al. (2022) for recent evidence and a summary.

The remainder of the paper is organized as follows. Section 2 describes the context of the Covid testing program. Section 3 develops the model and derives the test for social norms. Section 4 describes our data. Section 5 shows our empirical results, while section 6 briefly concludes.

2 The context

The purpose of the University of Cambridge Asymptomatic Covid-19 Screening Programme (henceforth, the Programme) was to identify students with *asymptomatic* infection, so they could avoid infecting other people. The Programme was in place for four academic terms, spanning two academic years.[9]

A cornerstone of the Programme and key to our statistical analysis is the concept of a 'student household' and we begin with a description of that (Section 2.1). This is followed by a description of the Programme itself (Section 2.2).

2.1 The student households

All University of Cambridge students, whether undergraduates or postgraduates (Master's or PhDs), belong to one of its 31 Colleges. Among other things, the Colleges provide accommodation to roughly two-thirds of all Cambridge students, or about 16,000 students a year. Students have their own room, but often share facilities – such as a kitchen and/or a bathroom – with others.

Most Colleges have both historical and more modern buildings on their sites, and some own terraced houses or blocks of flats in the city that are also used for student accommodation. Students are allocated to rooms in October before the start of each academic year for one year. In a 'normal' year, the allocation of rooms to returning students varies by College, with the majority of Colleges using some form of a ballot. In some Colleges, returning students can submit requests as a group of friends who would like to live together, particularly when the accommodation consists of self-contained houses or flats. In contrast, new incoming freshmen (all first year undergraduates and some new postgraduates) are allocated to rooms by college

9. At the University of Cambridge, the academic year is divided into three terms. Michaelmas (first) term is between October and December, Lent (second) term is between January and March, and Easter (third and final) term is between April and June.

administrators depending on where there is space.[10] So, arguably, freshmen allocations to rooms are near-random, at least conditional on observables.

The allocation of rooms during the academic years affected by the Covid-19 pandemic was done along similar lines, but with important adjustments. Before the start of the 2020-2021 academic year, the University, to comply with national Covid-19 regulations, defined what the concept of a 'household' would mean for the student population of the University. It settled on 'a group of rooms that shared facilities such as a kitchen and/or bathroom'.[11] This meant that, under national policy, if the U.K. went into lockdown, then the members of the household could continue to socialise with each other, but not with members of other households, or if any member of a household tested positive for Covid-19, all members would have to self-isolate for a prescribed period. The resulting 4,000 households within the University ranged in size from 1 to 20 students (though those with 10 or more were uncommon). The mean household size is just under six students and about 15% of the households had only one student living in them. Figure 1 shows the distribution of household sizes in October 2020.

Given the unique social importance of a household during the pandemic, a number of Colleges relaxed their normal balloting process for returning students and allowed more students to express a preference for being allocated to the same household as their friends. In response to a survey of college COVID-19 policies that we conducted in July 2021, 22 out of 28 the Colleges that responded reported that they continued to use ballots; 10 of these also allowed students to request to be in a household with friends. For our analysis, this means that although the prevalent method of room allocation for freshmen is nearly random, for returning students, there was some scope for self-selection. This is something we carefully address in the statistical analysis.

Under normal circumstances, Cambridge students are required to reside in Cambridge during term time, and this requirement is strictly enforced. However, during the second and third term of the 2020-21 academic year, the University made an exception due to the pandemic, allowing students to study away from Cambridge and return at a time of their choosing. This, among other things, meant that some students who decided to return to Cambridge found themselves in a 'singleton' household because all their housemates had decided not to

10. Although incoming freshmen often express preferences with regard to en-suite bathroom or price band on their applications, college administrators may not satisfy them.

11. An integral part the UK government's Covid-19 restrictions and guidance was the concept of a 'household'. In May 2020, a 'household' was defined as a group of people living together in the same home who shared facilities such as a kitchen, bathroom, or living space. Members of a household were considered a single unit for social distancing and isolation purposes.

Figure 1: Distribution of the size of the households within the 31 Cambridge Colleges in October 2020

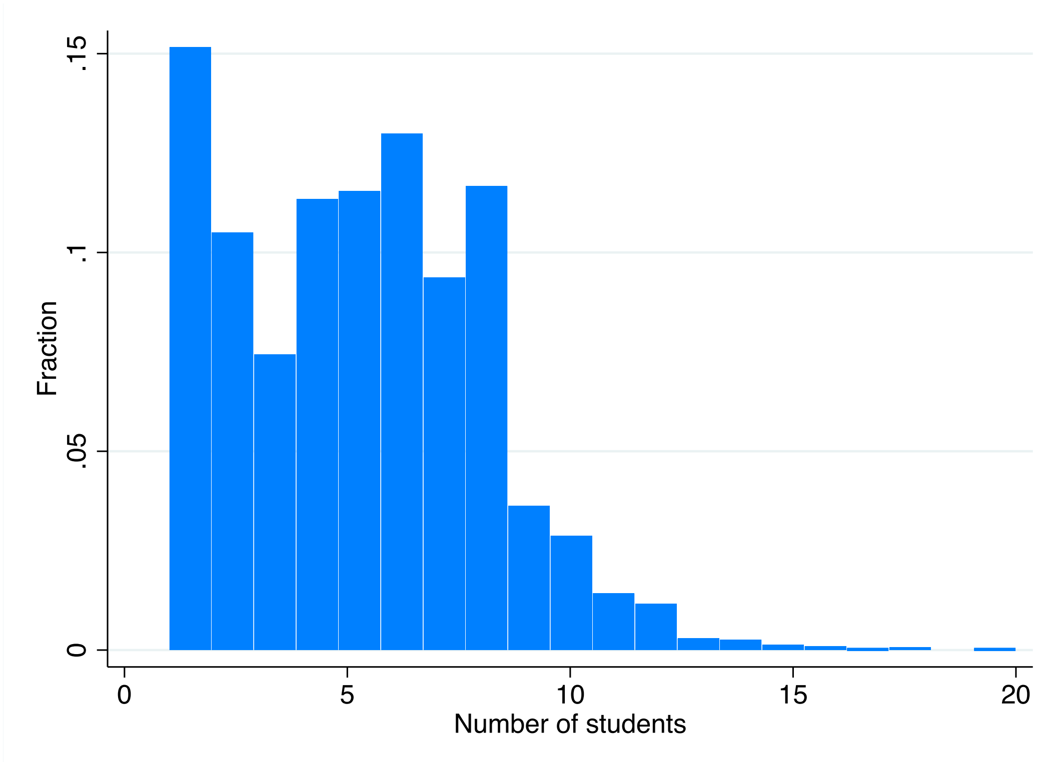

return.[12] Households with just one student could, therefore, be created by coincidence or by the architecture of a College that made it natural for the administrators who designed the household structure within a College to denote particular rooms as their own household. In our subsequent analysis, we exploit the variation created by some students living alone while others live in households with others.

2.2 The Programme

2.2.1 The design

Soon after the Covid-19 outbreak in early 2020, it became clear that humans could experience and transmit the virus without symptoms (Matheson et al. 2021), and that young adults were likely to be asymptomatic (University of Cambridge 2020). Hence, there was a serious concern that universities would become 'conduits for transmission due to extensive social networks of young adults, many of whom live communally, and in-person teaching of large

12. On average 45% of all students were present in the second term and 83% in the third.

groups' (Aggarwal, Warne, et al. 2022).

The University of Cambridge Asymptomatic Covid-19 Screening Programme was created to address this by organising weekly testing among its students. The Programme, which tested 10,000 students per week at its peak, was set up from scratch in a matter of weeks for the start of the 2020-2021 academic year in October 2020 (University of Cambridge 2022). It ran continuously until January 2022, so for all three terms of the academic year 2020-2021 and the first term of the academic year 2021-2022.

All Cambridge students – undergraduate as well as postgraduate students – living in college accommodation (on site or in external houses or flats) were eligible for the Programme. This included nearly 16,000 students per academic year (or 23,000 separate individuals across the two academic years). This is two-thirds of Cambridge's student body. At the beginning of each of the two academic years, students were asked to *voluntarily* sign up to the Programme.[13] The sign-up rate was high, nearly 85%. Those who signed up received a weekly invitation to take the test.

Every week, each household received an envelope with test swabs for all the students in that household who had signed up and *one* test tube. The envelope had to be picked up and later the test tube (with the used swabs in it) had to be dropped off by one of the students in the household at a designated collection point, of which there would be one or two within the perimeter of each of the 31 Colleges.

The Programme used pooled testing in order to reduce running costs. First, a student would use the individual swab to take a sample from their nose and throat. Second, all students in the household, who had swabbed, would put their swabs into one test tube (i.e., pooled them) before one of the students would return the tube to the collection point (University of Cambridge 2020).[14]

The Programme instructed the students to put the test tube in a shared area (kitchen or bathroom), and encouraged them to agree on the time when testing would happen in their household (University of Cambridge 2020). It is, therefore, very likely that the students observed each other's testing decisions, particularly as the households were small, with six

13. Students who had signed up at the beginning of the academic year 2020-2021 automatically stayed in the Programme for the next academic year (2021-2022), as long as they did not opt out or had graduated. Students, who had not signed up initially when the Programme was launched or were new, could, however, sign up at the beginning of the new academic year (2021-2022).

14. Apart from a handful of exceptions, the 'pool' – i.e., the group that used the same test tube – was the household.

students on average.

If the result of the household (pooled) test was positive, then all members of the household had to self-isolate, initially for 14 and subsequently for 10 days.[15,16]

Participation in the Programme was strictly voluntary. Access to individual participation information was highly restricted and was never shared with the students' colleges. The university made it clear that there would be no sanctions for not participating (University of Cambridge 2020) and under the rules of the Programme, neither the University nor the Colleges were allowed to use incentives to encourage testing or to mandate testing. Instead, the recruitment material focused on the social benefits of individual participation. The official website of the Programme opened with the question 'Why should I take part?' and the answer was 'in order to help the others and the community'. Most official communications from the Programme were accompanied by the phrase 'Thank you for helping us keep Cambridge safe!' and contained other reminders to the students about the social benefit of active participation.

2.2.2 The incentives

To understand the students incentives to take part and test and why the actions taken by the participants in the Programme are good examples of prosocial actions, we highlight the costs and benefits from testing seen from the perspective of the students. Our model, presented in Section 3, builds on these considerations.

As already mentioned, the *raison d'être* of the Programme was the large expected social public health benefit.[17] It is much less clear whether there were significant private benefits to the students from taking part; in fact, given that the students were likely to be asymptomatic even if infected, private benefits were arguably quite low.[18]This conjecture is corroborated by a student survey that we conducted in March 2021, where the surveyed students rated reasons related to private benefits as relatively unimportant for their decision to test, instead highlighting social benefits and concerns about the opinions of others in their households as

15. In the U.K., the self-isolation rules were set nationally, changing from 14 to 10 days at the end of 2020.
16. Students belonging to a household that had returned a positive Covid-19 test tube were subsequently tested individually.
17. The Cambridge medics who designed and ran the Programme have estimated that the Programme was a major contributor to the comparatively low Covid-19 infection rates in Cambridge (Aggarwal, Warne, et al. 2022).
18. Some private benefits could emerge around vacation times, when many students would be planning to go home to see relatives. However, students traveling abroad had access to tests outside the Programme.

key reasons. In contrast, the private costs of testing were not negligible: testing risked self-isolation and was associated with considerable physical discomfort and hassle of remembering to test by the weekly college-specific deadline, picking up the test package and later returning the test tube.

Thus, with negligible private benefits and significant private costs, we would, generally, not expect students left to themselves to take the test, unless they had some degree of prosocial preferences.[19] These arguments strongly support our working hypothesis that we can treat taking the test as an example of a prosocial action.

Additionally, given the small household size, it seems plausible that students' payoff from testing may depend on what others in their household are doing, which would create interdependency in testing decisions. This can arise for two main reasons: social norms and externalities[20].

First, given that households are small and testing is observable, social norms may form, creating an incentive for students to conform to the behaviour of their housemates. This is borne out in our survey of students, where many reported that their testing decisions were influenced by concerns about the actions or perceptions of others in their household with respect to Covid-19 testing. Such desire to adhere to social norms in the household tend to induce strategic complementarity in testing decisions (e.g., Granovetter 1978). Furthermore, different norms may emerge. On the one hand, some students believe that testing is 'the right thing to do' and this could help develop a pro-testing norm. On the other hand, some students may develop an anti-testing norm, where a deviation would be met with the disapproval from the housemates as the testing person risks bring quarantine upon everyone in the household. In sum, some households may develop a norm in which members are expected to test, whilst the opposite norm, that of not testing, may emerge in others.

19. The following exchange between two prominent economists (Steven Levitt and Paul Romer) about mass asymptomatic Covid-19 testing makes the same point: "I think a question that you run into is: well, why would anyone test? It's not that much fun to have that thing stuck up your nose. And Paul Romer, who I've interviewed for People I (Mostly) Admire, his solution involved testing every single American once a week. And you'd have to go to Walgreens or something and tip your nose back and have somebody, 52 times, stick this thing in your nose when you have no reason to think you're sick. And Paul just thought, 'Well, people will do it because it's the right thing to do.' And my thought was, 'You do that to me once, the chance that I'm ever going to show up in Walgreens again is essentially zero'." (Levitt 2020)

20. Another possibility is sharing of private information, say about the benefits of testing. However, given considerable publicity by the Programme as well as continued discussions in the media, we do not feel there was much scope for private information. Indeed, the empirical results we find are hard to reconcile with the information story.

[21]

Second, the private marginal cost of testing may be lower if others are testing too because others can pick up and drop off the kit and tube. Moreover, the probability of having to self-isolate may be less affected by a student's decision to test if others in the household are testing already. These cost externalities, by lowering the individual cost of testing when other household members also test, would create strategic complementarity in testing decisions[22].

Both of these reasons would generate similar peer effects in testing decisions. Our aim is to go beyond the peer effect estimates and to distinguish between these two reasons. This is empirically challenging, and, in the next section, we develop a simple theoretical model that makes it clear how we can overcome this challenge using the data generated by the Programme.

3 Model

We develop a model in which students randomly assigned to households decide to test or not. We consider three scenarios, indexed by $k \in \{I, E, S\}$: (i) a student's payoff of testing is independent of the actions of others in the household ($k = I$, section 3.2), (ii) household members exert positive payoff externalities on each other ($k = E$, section 3.3), and (S) students get a benefit from conforming to a social norm as embodied by the behaviour of others in their household ($k = S$, section 3.4). We highlight differences in incentives in E and S, and derive implications for the observable distribution of people testing across households relative to I. This allows us to construct an empirical test for social norms (section 3.5).

3.1 Basic setup

We consider an environment with a finite number of students, indexed by i. Students are randomly assigned to households, indexed by h, with H_h denoting the set of students assigned to household h. The number of students in household h is N_h. Each student $i \in H_h$ decides

21. We use the term 'social norm' to refer to a shared understanding of how the group should behave that may induce a student to take an action which, in absence of the norm, would not be privately optimal (see Elster 1989 among others). This is in contrast with the sometimes used definition of social norms as equilibrium selection rules in coordination games (e.g. Goyal 2023). See Bicchieri et al 2023 for a discussion of the two definitions.

22. Negative externalities are less plausible in our context, but we discuss this possibility in appendix C.

Figure 2: Household-level equilibrium outcomes with positive externalities relative to independent decisions

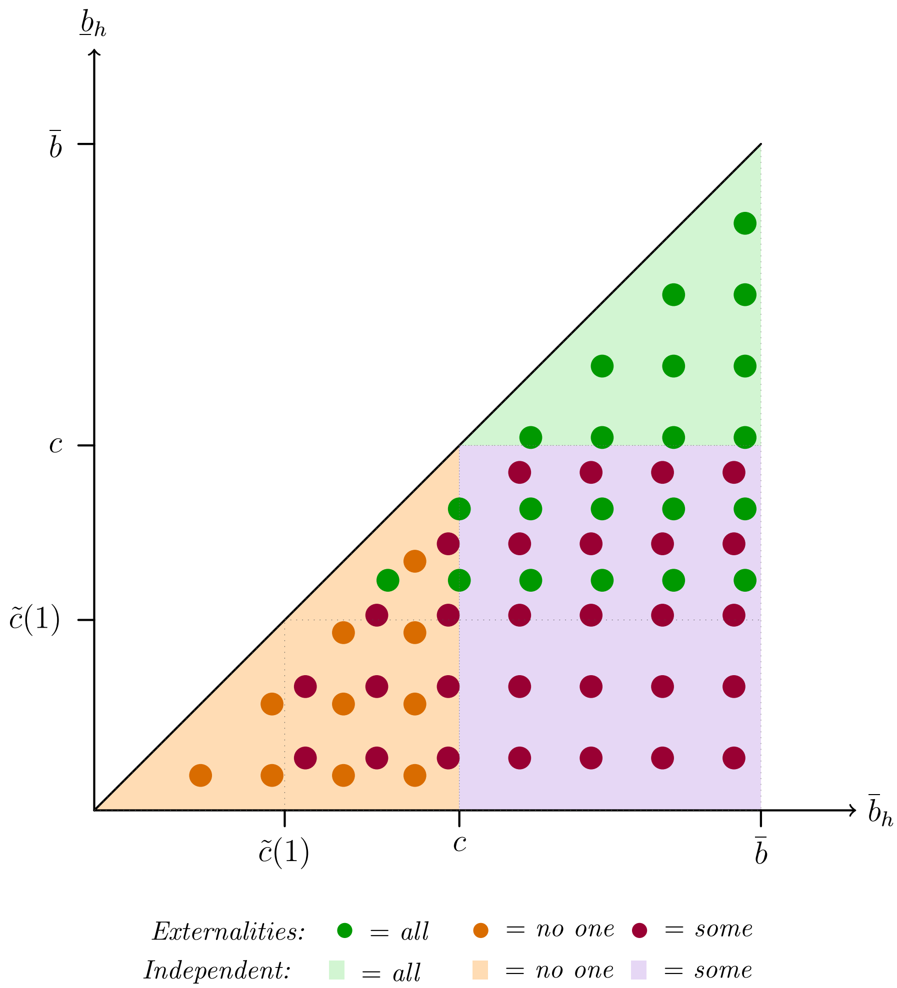

whether or not to test. We denote this decision by $t_h^i = \{0, 1\}$ and normalise the payoff of not testing to zero, i.e. $u_i(t_h^i = 0) = 0$ for all i. We summarise household-level outcomes by the number of students testing in household h (n_h) and focus on three categories of outcomes, indexed by $j \in \{A, N, M\}$: all test ($j = A$) with $n_h = N_h$, no-one tests ($j = N$) with $n_h = 0$, and a mixture with some students testing ($j = M$) so that $0 < n_h < N_h$.

All students have the same baseline cost of testing, $c > 0$, related to discomfort, hassle, and the expected cost of self-isolation, but differ in terms of benefits (b_i). In the student population, benefits are distributed on the interval $[\underline{b}, \bar{b}]$, drawn from the cumulative distribution function $F(b_i)$. We interpret the benefit as largely social, in the sense that students differ in how much they value the social benefits of testing, although this interpretation is not essential for the results of the model and there may also be a private component to

Figure 3: Household-level equilibrium outcomes with social norms relative to independent decisions

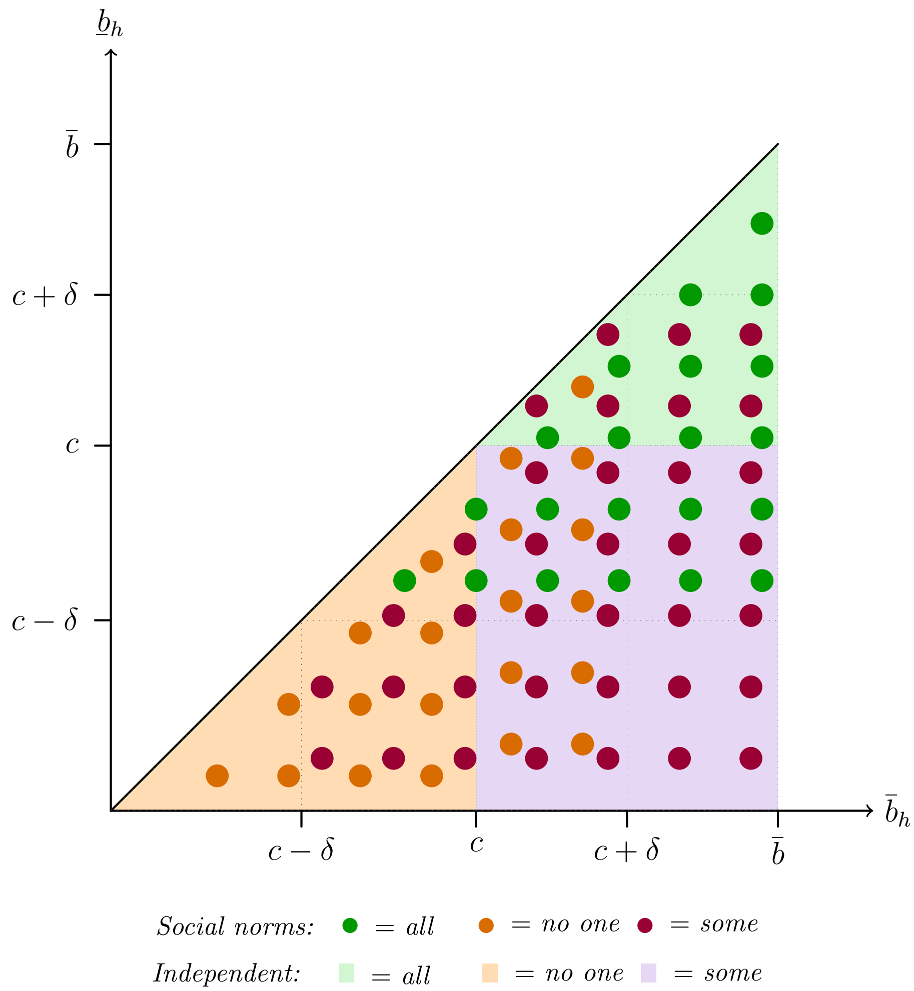

the benefit. Assuming that $\underline{b} < c < \bar{b}$, heterogeneity in benefits implies that, in isolation, some students will choose to test while others will not. Random assignment of students to households induces different (discrete) distributions of benefits across households: some households may be populated by students with high benefits, others by students with low benefits, and yet others by a mix of both. We characterise each household h by the member with the lowest benefit ($\underline{b}_h$) and the highest benefit ($\bar{b}_h$), so that benefits within the household are distributed on the interval $[\underline{b}_h, \bar{b}_h]$. These bounds are sufficient to determine which of the three outcome categories ($j \in \{A, N, M\}$) a household can fall into.

The payoff from testing for a student living in household h may depend on the testing decisions of others in the household (the relevant peer group). We model the testing decision within a household as a full-information, one-shot, simultaneous-move game, and solve for

Nash equilibrium values of n_h under three scenarios which we index by $k \in \{I, E, S\}$. Recall that the three scenarios are: independent payoffs ($k = I$), positive externalities ($k = E$), and social norms ($k = S$). The scenario with independent payoffs is the benchmark to which outcomes in the two other scenarios are compared.

3.2 Benchmark: Independent payoff ($k = I$)

In this benchmark scenario, we assume that there are neither household-level payoff externalities nor social norms. Students in a given household therefore decide whether or not to test based solely on a comparison between the baseline cost and their own benefit of testing, and test if and only if $u_i(t_h^i = 1) = b_i - c \geq u_i(t_h^i = 0) = 0$. The number of students testing in a given household (n_h) thus depends simply on how many household members satisfy $b_i \geq c$. The Nash equilibrium household-level outcomes are summarised in Lemma 1.

Lemma 1 (Benchmark) *The Nash equilibrium household-level outcome for household h of the simultaneous-move game with independent payoffs can be characterised as follows:*

- *All students in H_h test ($j = A$): $n_h^* = N_h$ if and only if $\underline{b}_h \geq c$.*
- *No student in H_h tests ($j = N$): $n_h^* = 0$ if and only if $\bar{b}_h < c$.*
- *Some students in H_h test ($j = M$): $n_h^* \in \{1, \dots, N_h - 1\}$ if and only if $\underline{b}_h < c$ and $\bar{b}_h > c$.*[23]

The lemma provides a benchmark for what the distribution of household-level outcomes would look like in the absence of any payoff interdependencies. Figures 2 and 3 visualise these outcomes. For now, ignore the circles. On the x-axis, we plot the maximum benefit of a student in a given household h ($\bar{b}_h$), and on the y-axis we plot the minimum benefit ($\underline{b}_h$). Each household is therefore located somewhere below the 45-degree line, depending on the distribution of benefits within it.

In the two figures, households located in the light green area are those in which all students test, because even the student with the minimum benefit prefers to do so. By contrast, in

23. This condition implies that some students will test, but not how many exactly. To pin this down, information about the particular distribution of benefits within the household is required. For our purposes, all that matters is that some students test, not the exact number. This remark applies to the following lemmas as well.

the light orange area, no one tests because it is not in the interest of even the student with the maximum benefit in that household to test. In the light purple area, some students test: the student with the highest benefit (and possibly some others as well) chooses to test, while the student with the lowest benefit (possibly along with others) does not.

The shares of households assigned to each of the three outcome categories, denoted S_j^I for $j \in \{A, N, M\}$, characterise the equilibrium distribution of outcomes across households. Under independent payoffs, these shares coincide exactly with the proportion of households for which the upper and lower bounds on the benefit size lie within the corresponding coloured regions in the two figures. The circles superimposed on the figures illustrate what happens when payoff interdependencies are introduced, in the form of positive externalities or social norms. We next turn to the analysis of these scenarios.

3.3 Externalities ($k = E$)

In this scenario, a student's decision to test lowers the cost of testing for others in the household; that is, the decision is associated with a positive externality.[24] Such externalities can arise for two reasons. First, when others in a household test, organisational and logistical hassle can be shared within the household (for example, it requires only one student to pick up or drop off the test kit for the entire household). Second, since students must self-isolate if at least one fellow student in the household tests positive, and infection risks among students living in close proximity are positively correlated, the marginal increase in the probability of self-isolation from a student's decision to test is lower when others in the household are testing as well.

To be concrete, we assume that a student's cost of testing decreases in the share of other students in their household who test, $\bar{t}_h^{-i} \equiv \frac{n_h^{-i}}{N_h - 1}$, where n_h^{-i} is the number of students in the household other than i who test. The cost of testing for a student in household h therefore becomes $\tilde{c}(\bar{t}_h^{-i})$, where $\frac{\partial \tilde{c}}{\partial \bar{t}_h^{-i}} < 0$, $\tilde{c}(0) = c$, and $\tilde{c}(1) \geq 0$.[25] The utility from testing for student

24. In general, there could also be negative household-level externalities if, for example, a student's percieved social benefit from testing is lower when others in the household test. Negative externalities would make testing decisions within a household strategic substitutes. We show in Section 5.1 that this is inconsistent with our data and, for that reason, we focus on the case with positive externalities in the text. For completeness, we show in Appendix C that the test we develop for distinguishing between externalities and social norms also works in the presence of negative externalities.

25. The following modifications do not affect our main results: (i) modelling the self-isolation externality separately, or (ii) assuming that one person bears all organisational costs, so that $\tilde{c}$ exhibits a downward jump only at $n_h^{-i} = 1$.

$i \in H_h$ is $u_i(t_h^i = 1) = b_i - \tilde{c}(\bar{t}_h^{-i})$. It is clear that testing decisions within a household are now interdependent and constitute strategic complements: if some students in a household test, it becomes more likely that others will also test because the cost of doing so is lower. The Nash equilibrium household-level outcomes are summarised in Lemma 2.

Lemma 2 (Positive externalities) *The Nash equilibrium household-level outcome for household h of the simultaneous-move game with positive externalities can be characterised as follows:*

- *All students in H_h test ($j = A$): $n_h^{**} = N_h$ if and only if $\underline{b}_h \geq \tilde{c}(1)$.*
- *No student in H_h tests ($j = N$): $n_h^{**} = 0$ if and only if $\bar{b}_h < c$.*
- *Some students in H_h test ($j = M$): $n_h^{**} \in \{1, \ldots, N_h - 1\}$ if and only if $\underline{b}_h < c$ and $\bar{b}_h \geq \tilde{c}(1)$.*

Proof. Given the fraction of other students in household h who test, student $i \in H_h$ will test if and only if $b_i - \tilde{c}(\bar{t}_h^{-i}) \geq 0$. If all other students in household h test, student i will also test if $b_i - \tilde{c}(1) \geq 0$. This condition holds for all students $i \in H_h$ if it holds for the student with the minimum benefit, that is, if $\underline{b}_h \geq \tilde{c}(1)$. In this case, it is a best response for all students in the household to test when others are testing. If no other student in the household tests, student i decides not to test when $b_i - \tilde{c}(0) < 0$. Since $\tilde{c}(0) = c$, it is a best response for all students in a household not to test when everyone else also abstains if the student with the maximum benefit chooses not to test, that is, if $\bar{b}_h < c$. In households with $\underline{b}_h < c$ and $\bar{b}_h \geq \tilde{c}(1)$, some students—but not all—will test ■

Lemma 2 has two important implications. First, compared to the scenario with independent payoffs, households with $c > \bar{b}_h > \underline{b}_h > \tilde{c}(1)$, which under independence sustain a unique 'no one tests' equilibrium, can, because of the positive externality, sustain an equilibrium in which everyone tests. Similarly, households with $\bar{b}_h > c > \underline{b}_h > \tilde{c}(1)$, which under independence would sustain equilibria with some testing, can also sustain an equilibrium with everyone testing when positive externalities are present. This is illustrated in Figure 2, where the dark green circles indicate the range of parameter values for which households can sustain equilibria with all students testing. We see that the green circles expand into the orange and purple areas. For the same reason, the range of parameter values for which households can sustain equilibria with some testing—illustrated by the purple circles—also expands into the orange area. Importantly, the range of parameters for which the 'no one tests' equilibrium can be sustained is the same as under independence. The intuition is that a positive externality cannot induce a student who would otherwise choose to test to stop testing. This implies that a household with some testing *cannot* be transformed into a household with no testing as a consequence of positive payoff externalities.

Second, positive externalities render testing decisions strategic complements within households. As a result, there now exist households in which multiple equilibria may be played, illustrated by areas with circles of multiple colours. An implication of this is that we cannot say exactly how large a share of households will fall into which outcome category without imposing an equilibrium selection rule. Importantly, however, we can put bounds on the shares *relative* to the scenario with independent payoffs. This is summarized in the following proposition.

Proposition 3 *The distribution of equilibrium outcomes with positive externalities ($k = E$) compared to the distribution with independent payoffs ($k = I$) must satisfy*

$$S_N^E \leq S_N^I$$
$$S_A^E \geq S_A^I$$

Proof. The share of households in which no student tests in the presence of positive externalities, S_N^E, cannot exceed the corresponding share under independent payoffs, S_N^I, because no student who chooses to test without the externality would prefer not to test once the externality is introduced. This share may, however, be smaller if some households with parameters in the range $\tilde{c}(1) < \bar{b}_h < c$ and $\tilde{c}(1) < \underline{b}_h < c$ switch from the equilibrium in which no student tests to an equilibrium in which some or all students test. Similarly, the share of households in which all students test in the presence of positive externalities, S_A^E, cannot be smaller than the corresponding share under independent payoffs, S_A^I, for the same reason: no student who tests without the externality would prefer not to test with it. Consequently, households in which all students test under independent payoffs will continue to play that equilibrium when positive externalities are present. Moreover, some households with parameters in the range $\tilde{c}(1) < \underline{b}_h < c$, which under independent payoffs play an equilibrium with partial or no testing, may switch to the equilibrium in which all students test once the externality is introduced ■

The proposition shows that with positive externalities the share of households in which no one tests will be weakly smaller than under independence, while the share of households in which everyone tests will be (weakly) larger. The intuition is that no student who tests without the externality would prefer not to test with it. Taken together, this means that, relative to students deciding independently, total testing across all students cannot fall, and may plausibly increase, when there are positive externalities.

3.4 Social norms ($k = S$)

In this scenario, we assume that students' testing decisions are influenced by household-level social norms. A student may conform to the prevailing norm in the household by emulating the testing choices of peers, even if doing so is not strictly in the student's private interest, because following the norm makes the student feel better (see, e.g., Bicchieri, Muldoon, and Sontuoso 2023; Elster 1989). Different households may develop different norms, i.e. norms that encourage testing and norms that discourage it. Following Blume et al. (2015), we model norms by introducing a 'conformity' component into the utility function that penalises a student if the student's testing decision deviates from the average testing decision of household peers.[26] Formally, the utility that student i in household h derives from testing is $v_i(t_h^i = 1) = b_i - c - \delta(1 - \bar{t}_h^{-i})^2$, and from not testing is $v_i(t_h^i = 0) = -\delta(-\bar{t}_h^{-i})^2$, where $\delta > 0$ is the weight placed on conformity and $\bar{t}_h^{-i} = \frac{n_h^{-i}}{N_h - 1}$ is the average testing decision of others in the household, excluding student i. We observe that student i's utility from testing relative to not testing is $\Delta v_i = b_i - c - \delta(1 - 2\bar{t}_h^{-i})$, and that this expression is increasing in the share of others in the household who test. Hence, testing decisions within a household are strategic complements, as in the externalities case (see Section 3.3). This is an example of the well-known non-identifiability result in the peer effects literature: externalities and social norms cannot, in general, be distinguished by deriving or estimating best-response functions (Brock and Durlauf 2001; Boucher et al. 2014; Boucher et al. 2024). The Nash equilibrium household-level outcomes are summarised in Lemma 4.

Lemma 4 (Social norms) *The Nash equilibrium household-level outcome for household h of the simultaneous-move game with social norms can be characterised as follows:*

- *All students in H_h test ($j = A$): $n_h^{***} = N_h$ if and only if $\underline{b}_h \geq c - \delta$.*

26. A similar parametrisation is used by Manski (1993), Brock and Durlauf (2001), and many others, whilst Bernheim (1994) provides micro foundations for it. Boucher et al. (2024) experiment with using other moments (mode, maximum, minimum) of the distribution of peer decisions as the reference point. In our context, it makes no substantive difference which reference point is used.

- *No student in H_h tests ($j = N$): $n_h^{***} = 0$ if and only if $\bar{b}_h < c + \delta$.*
- *Some students in H_h test ($j = M$): $n_h^{***} \in \{1, \ldots, N_h - 1\}$ if and only if $\underline{b}_h < c + \delta$ and $\bar{b}_h > c - \delta$.*

Proof. Given the fraction of other students in household h who test, student $i \in H_h$ will test if and only if $\Delta v_i = b_i - c - \delta(1 - 2\bar{t}_h^{-i}) \geq 0$. If all students other than i test in household h, student i will also test if $\Delta v_i = b_i - c - \delta(1 - 2) = b_i - c + \delta \geq 0$. This condition holds for all students $i \in H_h$ if it holds for the student with the minimum benefit, that is, if $\underline{b}_h \geq c - \delta$. In this case, it is a best response for all students in the household to test when others are testing. If no other student in the household tests, student i decides not to test when $b_i - c - \delta(0 - 2) = b_i - c - \delta < 0$. It is a best response for all students in a household not to test when everyone else also abstains if the student with the maximum benefit chooses not to test, that is, if $\bar{b}_h < c + \delta$.

To derive the condition under which some but not all students test in household h, note that Δv_i must be positive for at least one student and negative for at least one student in that household. This requires that

$$\bar{b}_h > c + \delta(1 - 2\bar{t}_h^{-i}) \quad \text{for all } \bar{t}_h^{-i}, \tag{1}$$

and

$$\underline{b}_h < c + \delta(1 - 2\bar{t}_h^{-i}) \quad \text{for all } \bar{t}_h^{-i}. \tag{2}$$

Condition 1 is satisfied if it holds for $\bar{t}_h^{-i} = 1$, which implies $\bar{b}_h > c - \delta$. Condition 2 is satisfied if it holds for $\bar{t}_h^{-i} = 0$, which implies $\underline{b}_h < c + \delta$. ∎

Lemma 4 has two important implications. First, compared to the case with independent payoffs, all three types of equilibria – everyone tests ($j = A$), no one tests ($j = N$), and some students test ($j = M$) – become possible for a larger set of households. This is illustrated in Figure 3, where the dark green circles indicate that an equilibrium in which all students test becomes possible in households that under independence would either have no one testing (located in part of the light orange area) or some testing (located in part of the light purple area). The range of parameters for which some testing occurs also expands. Most importantly, however, the range of parameters for which an equilibrium with no one testing can be sustained also expands relative to the case with independent payoffs. This contrasts with the scenario with externalities. The reason is that a household that, in the absence of norms, would sustain equilibria with either some or all students testing can, in the presence of norms, sustain an equilibrium in which no one tests because this becomes the household norm.

Second, social norms make testing decisions strategic complements and induce multiple equilibria for all households located in the regions of Figure 3 indicated by circles of different colours. One implication is that we cannot unambiguously predict how the shares of households that play each of the three types of equilibrium outcomes will change relative to the scenario with independent payoffs, as this depends on which equilibrium each household selects. Nor can we, as in the scenario with positive externalities, place bounds on these shares. What we can do is highlight a feasibility result in the following proposition.

Proposition 5 *It is feasible that the distribution of equilibrium outcomes with social norms ($k = S$), compared to the distribution with independent payoffs ($k = I$), satisfies*

$$S_N^S > S_N^I,$$
$$S_A^S > S_A^I.$$

Proof. The shares of households in which no student tests or all students test in the presence of social norms, S_N^S and S_A^S, can exceed the corresponding shares under independent payoffs, S_N^I and S_A^I. This is because, depending on the prevailing household norm, a student who chooses to test in the absence of a norm may prefer not to test when a norm is present, while a student who chooses not to test without a norm may be encouraged to test with one.

Such an outcome can arise, for example, if all households that, under independent payoffs, play either the 'no one tests' or the 'everyone tests' equilibrium continue to do so when social norms are present, while some households that play the 'some test' equilibrium in the parameter ranges $c < \bar{b}_h < \bar{b}$ and $0 < \underline{b}_h < c$ switch to either the 'all test' or the 'no one test' equilibrium ■

The proposition shows that, unlike in the scenario with positive externalities, it is possible for the share of households in which no one tests to *increase* relative to the case with independent payoffs. Consequently, total testing may fall—though this is not inevitable—if a norm of not testing emerges in a sufficiently large number of households.

3.5 Test for social norms

Propositions 3 and 5 show that the type of social influence operating within households — externalities or social norms — affects the distribution of equilibrium outcomes *relative* to the benchmark where the payoffs from testing are independent. This distinction allows us, despite the fact that both types of interdependency imply strategic complementarity and upward-sloping reaction functions, to develop an empirical test that differentiates between the two and detects the presence of social norms in a given environment.

To see how this works, suppose that the observed distribution of outcomes across households is $\{\hat{S}_A, \hat{S}_N, \hat{S}_M\}$, and that the corresponding (counterfactual) shares in the absence of any social influence are $\{\hat{S}^I_A, \hat{S}^I_N, \hat{S}^I_M\}$. The two propositions show that both externalities and social norms can increase the share of 'all test' households relative to the scenario with independent payoffs, but that only social norms can increase the share of 'no one tests' households. With externalities, this share either remains unchanged or declines. This theoretical insight forms the basis of our test:

Proposition 6 (The distribution test) *Suppose that the observed (empirical) distribution of household-level equilibrium outcomes is* $\{\hat{S}_A, \hat{S}_N, \hat{S}_M\}$*. If*

$$\hat{S}_N > \hat{S}^I_N, \tag{3}$$

then the data-generating process must have involved social norms.

Proposition 6 shows that an increase in the share of 'no one tests' households is a litmus test for the presence of social norms: it is a sufficient condition to establish that social norms are at work. In short, if social norms are sufficiently strong, we can overcome the non-identifiability problem and establish whether social norms and conformity are present. This does not rule out the possibility that externalities also play a role, but it does rule out the data being generated solely by that mechanism.

4 Data

Our dataset includes weekly information collected by the Cambridge's Asymptomatic Screening Programme, matched with data on student characteristics and household composition.

In total, the dataset comprises approximately 23,000 individual students living in over 4,000 households. The key outcome variables are students' decisions to join the Programme and, conditional on participation, their weekly decisions on whether to take a Covid-19 test. The Programme operated for the total of four terms: all three terms of the 2020–21 academic year, and the first term of 2021-2 academic year. It had not started collecting weekly student-level testing data until the second term of 2020-1. So these testing data are available for three terms: the second (Lent) and third (Easter) terms of 2020-1 and the first (Michaelmas) term of 2021-2, and they form the Programme-related part of our data set.

Table 1: Summary statistics for students eligible for participation in the Programme and selected Programme outcomes.

	2020-1 (2 terms)	2020-2 (1 term)
General characteristics:		
Number of students	15,623	16,095
Share of females	0.47	0.48
Share of white students	0.62	0.59
Share of UK students	0.67	0.67
Share of postgraduates	0.29	0.30
Seniority	1.99	2.05
Related to the Programme:		
Share living alone	0.03	0.05
Share ever signed up to the Programme	0.85	0.83
Present in a given week (prob.)	0.71	1.00
Tests in a given week (prob.)	0.45	0.27

Notes: The general characteristics refer to all students eligible for the Programme. The 'share of' variables range between 0 and 1. Seniority is defined as an undergraduate student's current year at Cambridge (fresher = 1, etc.). Undergraduate degrees at Cambridge are typically three years, but in a few cases four years. 'Present in a given week' is proxied by a 1-0 indicator of whether the student connected to the University Wi-Fi during a given week, averaged across weeks and students. 'Tests in a given week' is a 1-0 indicator of whether a student tested in a week when they were present and signed up for the Programme, averaged across such weeks and students. For some variables, the number of observations is slightly smaller than the number of students due to missing values.

Table 1 provides summary information on the characteristics of all students in the dataset (i.e., those eligible for the Programme) and on selected aggregated Programme outcomes. As shown in the table, the student body is highly international, ethically diverse, and gender-balanced. Undergraduates are more likely to live in college accommodation than postgraduates and are therefore more heavily represented in our data than in the overall student population. About 5% of students live alone, while the rest live in a household with at least one other person.

Nearly 85% of eligible students signed up for the Programme, with similar rates across the two academic years. Individual testing data available from the second term of the program imply the average testing rate of 45% in the 2020-21 academic year, dropping to 27% in the first term of 2021–22[27].

Table 2 reports summary statistics for the households. The average household size is nearly six students, with substantial variance (as illustrated by Figure 1). Households tend to be diverse in terms of gender, ethnicity, and country of origin. At the same time, postgraduate and undergraduate students are rarely assigned to the same household. There are no notable differences in these characteristics between the two academic years.

A significant proportion of households (around one-fifth) consists solely of first-year students (fresher households; columns 2 and 5). These households are broadly similar to other households, with only minor differences in the share of females and in ethnic and country of origin composition. Singleton households (households with only one student; columns 3 and 6) account for 15% of all households in the 2020–21 academic year and 23% in 2021–22. These households contain a higher share of postgraduates and male students and a lower share of UK and white students than multi-student households. During the two terms in the 2020-21 academic year, when the University allowed students not to return to Cambridge, students assigned to multi-student households had a 15-18% chance of finding themselves living in an 'accidental' singleton household.

At the household level, we highlight two stylised facts about students' testing behaviour that inform the formal analysis below. First, the last row of Table 2 reports, for each household type, the probability that a test tube is returned in any given week, indicating that at least one student in the household has tested. The table shows that students living alone are less likely to return a test tube than those living with others, indicating systematic differences in testing behaviour between singleton and multi-student households.

27. This is conditional on signing up for the Programme and being present in Cambridge.

Table 2: Household characteristics

	2020-1 (two terms)			2021-2 (one term)		
	Main	Fresher	Singleton	Main	Fresher	Singleton
	(1)	(2)	(3)	(4)	(5)	(6)
Number of households	1,931	674	461	1,986	575	764
Number of students	5.75	5.73	1.00	6.10	5.58	1.00
Share of females	0.45	0.51	0.39	0.47	0.50	0.42
Gender HHI	0.69	0.68	1.00	0.69	0.69	1.00
Share of white students	0.65	0.62	0.57	0.62	0.58	0.57
Ethnicity HHI	0.62	0.57	1.00	0.59	0.56	1.00
Share of UK students	0.66	0.68	0.43	0.66	0.66	0.54
Country HHI	0.63	0.66	1.00	0.63	0.66	1.00
Share of postgraduates	0.29	0.27	0.74	0.31	0.24	0.56
Post/undergrad HHI	0.95	0.99	1.00	0.96	0.99	1.00
Average seniority (years)	2.39	1.00	2.23	2.40	1.00	2.09
Seniority HHI	0.70	1.00	1.00	0.69	1.00	1.00
Accidental singleton (prob.)	0.15	0.18	0.00	0.00	0.00	0.00
Test tube returned (prob.)	0.76	0.67	0.61	0.47	0.50	0.28

Notes: Household groups: 'Multi' (columns 1 and 4) refers to multi-student households with at least two students and at least some students who are in their second year (i.e., not all-fresher households). 'Fresher' (columns 2 and 5) refers to households consisting only of first-year students. 'Singleton' (columns 3 and 6) refers to households consisting of one student only. Variables: 'Number of students' is the average number of students per household in the corresponding household group. 'Share of' variables range from 0 to 1. 'HHI' stands for Herfindahl-Hirschman Index, where 1 corresponds to one group only being present, and values below 1 correspond to more diverse households. The 'Country HHI' is calculated from data on the student's country of origin.'Average seniority (years)' is the average year at Cambridge, averaged across students in the household (i.e., freshers = 1, sophomores = 2, etc.). 'Accidental singleton (prob.)' is the probability of a student being left on their own in a multi-person household in a given week during the second or third term because the other students in the household decided to reside outside Cambridge. 'Test tube returned (prob.)' is the probability that a household returns the test tube in any given week.

Second, there is substantial heterogeneity in the number of students who test within households of similar size. Figure 4 provides an illustrative example. It shows the distribution of the number of students testing in households with four students[28]. The figure makes it clear that there are sizeable groups of multi-student households in which no one, or almost everyone, tested regularly, alongside a range of households in which only some students tested on a regular basis.

28. The figure is for Lent and Easter terms. In the Michaelmas term of 2021, the last term of the Programme, the distributions are similar, though shifted to the left as testing declined over the course of the pandemic (see Appendix B for all terms and household sizes).

Figure 4: The distribution of the number of students testing weekly in households of four people.

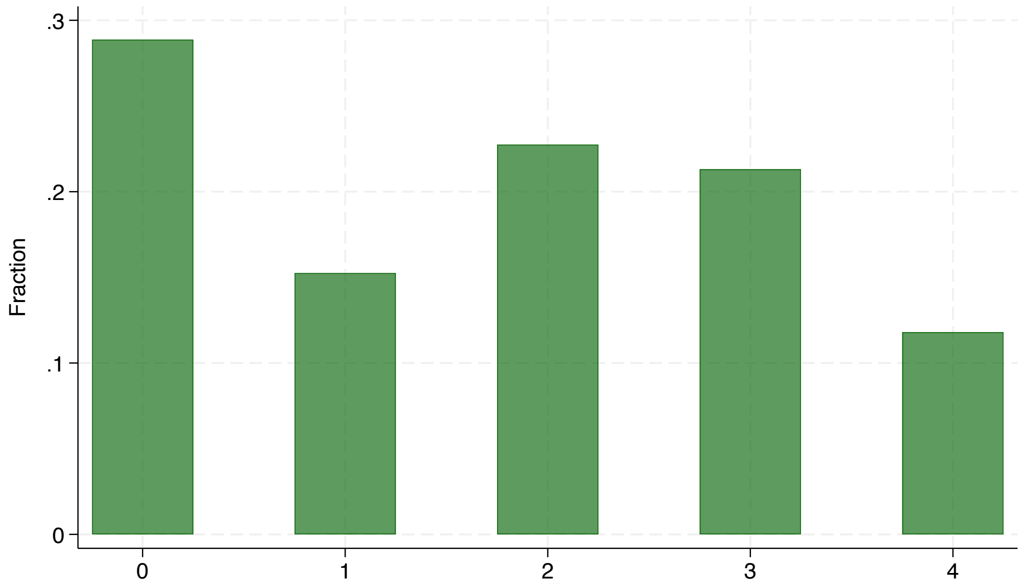

Notes: Lent and Easter terms 2021.

5 Empirical results

Both externalities and social norms predict that the decisions of students in the same household have causal positive influences on each other. Empirically, we, first, show that the data support this prediction by estimating peer effects, which are large and positive. Second, we separate the two mechanisms by implementing our distribution test, and show that the data point at the presence of social norms.

5.1 Peer effects

Students make two decisions: whether to sign up to the Programme, and then, having signed up, whether to test in any given week. We first estimate peer effects in testing, conditional on signing up, using the standard linear-in-means model:

$$\Pr(t_{ih\tau} = 1) = f(\alpha + \beta \bar{t}_{h\tau}^{-i} + \delta X_i + \eta_i + \mu_h + \nu_\tau + \varepsilon_{ih\tau}), \tag{4}$$

where $t_{iht} = 1$ is the testing decision of student i in household h and week τ, and $\bar{t}_{h\tau}^{-i}$ is the share of students in household h who decide to test in week τ, excluding student i. The vector X_i are observable student characteristics, η_i is the effect of unobservable student characteristics, μ_h is the effect of unobservable household characteristics, ν_τ are the aggregate weekly shocks, and $\varepsilon_{ih\tau}$ is the error term.

There are well-known challenges to identification of peer effects (Manski 1993; Brock and Durlauf 2001). They fall into three categories 1) Self-selection on unobservables related to testing – e.g., if the students choose their housemates based on their degree of prosociality, testing decisions will be positively correlated within households even in absence of peer influences; 2) Contextual effects, when unobserved characteristics of students simultaneously affect their own decisions and also the decisions of their housemates – e.g., if some students in a household are particularly vulnerable to Covid-19, they might test more often, and their housemates might test more often out of consideration for them; 3) Correlated effects, when some households have features that affect testing decisions of everyone in it – e.g., if a household is located a considerable distance away from the test collection point.

We address these challenges in the following ways:

1. Self-selection: As most students are allocated to rooms via ballots or through near random allocation by administrators, only a minority of households are subject to self-selection. Nevertheless, for robustness, we estimate all of our models on the administrator-allocated freshers as well as on the whole sample. As we will see, the estimates are very similar.

2. Contextual effects: As allocation to households is near random, the possibility of correlated effects is reduced. However, since households are small, they may have significant differences in composition by chance, and this can generate contextual effects. We exploit the panel structure of our data and, following Brock and Durlauf (2007), use student fixed effects to control for all unobserved student characteristics (η_i), thus eliminating contextual effects.

3. Correlated effects: Similarly, we control for household fixed effects to capture any time-invariant household characteristics that may affect testing (μ_h)[29], as well as week effects to capture any changes in aggregate variables that affect testing in general (ν_τ).

29. Household composition stays the same in any given academic year, but is different across years, which means student and household fixed effects can be estimated separately.

We summarise our results in Table 3.[30] We find strong positive peer effects: when everyone else in a student's household switches from not testing to testing, this student's probability of testing goes up by around 70 percentage points. Limiting the sample to freshers only does not change the results, which suggests that self-selection is not a major issue in our data.[31]

Table 3: Peer effects in testing

	All students		Freshers
	(1)	(2)	(3)
Testing rate among others in household	0.776***	0.623***	0.738***
	(0.01)	(0.01)	(0.02)
Household FE		✓	✓
Student FE		✓	✓
Week effect	✓	✓	✓
N	224,752	224,089	48,937

Notes: The dependent variable is a dummy that equals to 1 if the student takes the test. The estimator is OLS. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

The key identifying assumption in the regressions reported in Table 3 is that weekly shocks experienced by each student, $\varepsilon_{ih\tau}$, are uncorrelated across students in the same household. This may not be the case. To address this, we construct two robustness checks. First, since much of the typical student's social life happens at the college level, we include college-week effects that capture weekly shocks specific to one of the 31 colleges. Second, we use the household testing rate lagged one week ($\tau - 1$) to instrument for the household testing rate in week τ. The identifying assumption now is that whilst student-specific shocks may be correlated in the household in the same week, the shock to i in one week is not correlated to the shock to j in another week. This is a plausible assumption, particularly as we control for aggregate weekly shocks. The results of these robustness tests are summarised in Table 4. We continue to estimate positive, statistically significant and large peer effects.[32]

Before the students can take the weekly test, they have to sign up to the Programme. It is,

30. We estimate a linear probability model, for transparency and ease of interpretation.

31. Inclusion of fixed effects also does not change the results very much, suggesting that student and household characteristics do not play a big role in explaining the correlation between the students decisions.

32. Unsurprisingly, the IV coefficients are somewhat lower than those in the main estimates reported in Table 3. One reason is that IV regressions estimate the causal flow of peer effects from $j \neq i$ to i, but not from i to $j \neq i$, whilst the regressions reported in Table 3 include all of the mutual influences experienced in a Nash equilibrium. As our main goal is to establish whether there are causal peer effects, their direction and total magnitude, but not who is influencing whom, the presence of this 'reflection' problem in the main specification is not a particular concern.

Table 4: Peer effects in testing: Robustness

	College-week effects		IV using one-week lag	
	All students (1)	Freshers (2)	All students (3)	Freshers (4)
Testing rate among others in household	0.586***	0.648***	0.322***	0.632***
	(0.01)	(0.02)	(0.04)	(0.15)
Household FE	✓	✓	✓	✓
Student FE	✓	✓	✓	✓
Week effect			✓	✓
College-week effect	✓	✓		
N	224,502	48,937	196,277	42,760

Notes: Dependent variable is a dummy that equals to 1 if the student takes the test. Estimition is OLS. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

therefore, interesting to ask whether there are also peer effects in the decision to sign up (the extensive margin). Most students, once they sign up, stay signed up with the Programme (if they change their mind, they can simply stop testing.) Thus, overtime variation is limited, and we cannot use fixed effects, making it harder to rule out other sources of correlation between decision of housemates to sign up. Still, we can address self-selection through freshers, some of the contextual effects by controlling for student characteristics, and some of the correlated effects through year dummies. The results in Table 5 show positive, significant and reasonably large peer effects in signing up: When everyone else in the household joins the program, a student's sign up probability rises by about 40 percentage points.

Table 5: Peer effects in signing up

	All students (1)	Freshers (2)
Sign up rate among others in household	0.41***	0.38***
	(0.02)	(0.04)
N	29,614	6,832

Notes: Dependent variable is a dummy which equals to one if the student signs up to the Programme. Regressions include controls (gender, domestic student dummy, postgraduate student dummy, year at university, size of household, and academic year effects). Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

5.2 Social norms versus externalities

The positive household-level peer effects that we find confirm that testing decisions are strategic complements. This is consistent with the presence of both social norms and externalities. To distinguish between these mechanisms, we implement the distribution test developed in Section 3.

The test requires us to compare the distribution of the number of students testing across households with the distribution that would arise under independence. If social norms of both testing and not testing are present, then, relative to independence, we should observe an increase in *both* the share of households in which all students test and the share of households in which no one tests. Under externalities, again relative to independence, the share of households in which all students test will increase but, critically, the share of households in which no one tests will fall (see Proposition 6).

A key feature of our data is that some students (around 5%) live alone, and we observe their testing behaviour. To construct our independence benchmark, we proceed in three steps. First, we construct 'placebo' households by randomly allocating students who live alone to an imaginary household of a fixed size. Second, we calculate the total number of students testing in each 'placebo' household. Third, we compare the distribution of the number of students testing in actual households of the same size with the corresponding distribution from the placebo households, in which students behave independently.

Figure 5 illustrates this comparison for households with six students in the Lent and Easter terms 2021. The placebo households exhibit a distribution that is close to binomial, as expected under independent decision-making. Relative to this benchmark, the actual distribution shows both a higher share of households in which all students test and a higher share of households in which no one tests. This pattern is consistent with households exhibiting different social norms and is not consistent with externalities being the sole mechanism at play. The data display the same pattern across all terms and household sizes (see the graphs in Appendix B).

Although most students in our data cannot self-select into households, there is a concern that students who live alone are not a random subset of all students, which could bias our test. To address this concern, we construct alternative placebo households using students who ended up living alone in Lent and Easter 2021, when the University allowed students not to return to Cambridge after the Christmas vacation and to study remotely instead. These

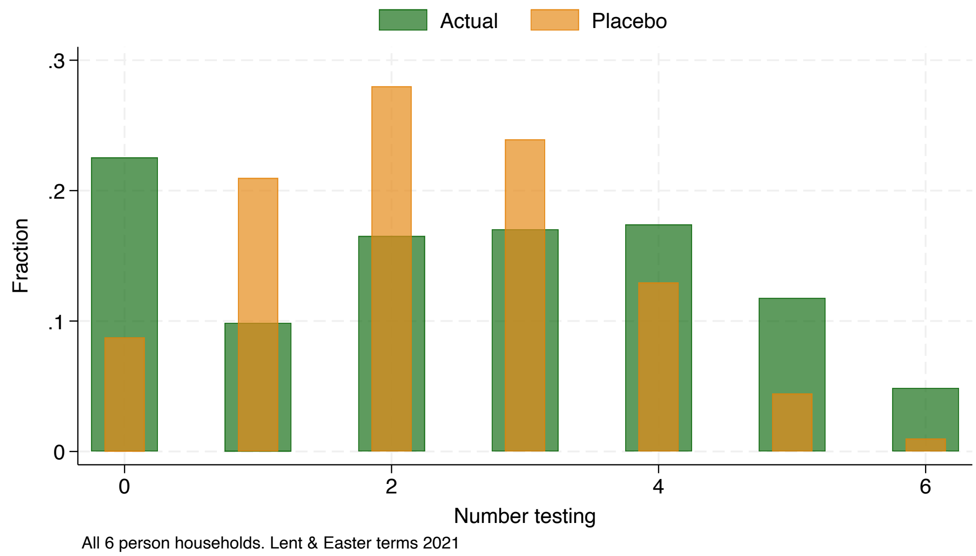

Figure 5: Number of students testing: actual versus placebo (independence benchmark) households

students had expected to live with others but found that they were the only student from their household to return. To further mitigate any residual self-selection, we also repeat the test using only freshers, in both placebo and actual households. The results for six student households are shown in Figure 6 and look very similar to those reported in Figure 5.

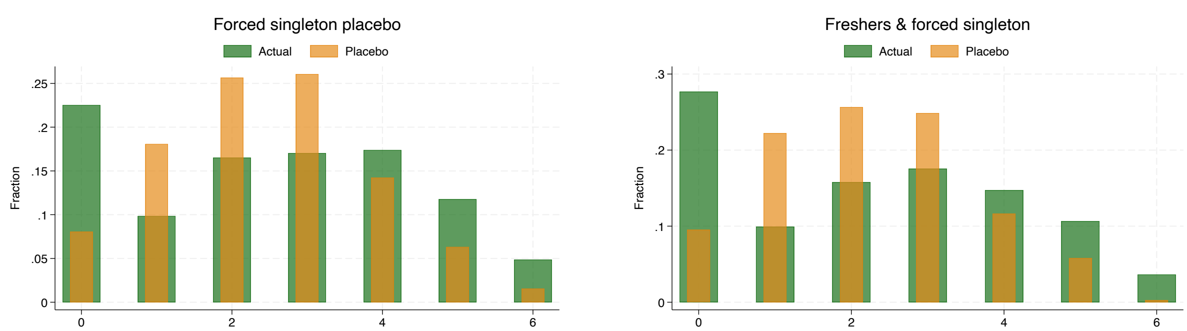

Figure 6: Number testing: Robustness

We now assess whether the differences that are clearly visible in the graphs are statistically significant. Table 6 summarises the results.[33] First, we compare the share of households

33. For consistency across columns, all results reported in Table 6 use data from two of the three terms covered in our dataset: Lent and Easter 2021. Only in these terms did the University allow students not to be physically present in Cambridge. Consequently, the 'forced' singleton placebo households in columns 3 and 4 can be constructed only for these terms. Results for the remaining term, Michaelmas 2021—the final term of the Programme—are reported in Appendix D and are qualitatively similar.

in which all students test in the actual households with that in the placebo households (column 1) and find that the former is larger, with the difference statistically significant at the 1% level. Second, we compare the share of households in which no one tests across the two groups (column 2). Again, we find that this share is higher in the actual households, with the difference significant at the 1% level. While the first finding is consistent with both social norms and externalities, the second is consistent only with social norms. We therefore conclude that students' testing decisions are influenced by social norms. The results are unchanged when we use 'forced' singleton households—students who ended up living alone because others did not return—to construct the placebo households (columns 3 and 4). Finally, the findings also remain the same when we restrict attention to freshers only (columns 5 and 6). Hence, our results are not driven by selection.

In conclusion, the testing data cannot be explained by externalities alone, but are consistent with conformity to social norms. Different households appear to develop different norms: some households adopt a norm of testing, while others adopt the opposite norm of not testing. This pattern is consistent with the model developed in Section 3 and with the broader theoretical literature on social norms, which predicts multiple equilibria and the emergence of different norms across peer groups.

Focusing on the estimates in column 3, the magnitudes imply that the share of households in which no one tests is 10 percentage points higher among actual households than among placebo households. This provides a quantitative estimate of the effect of a negative social norm on student behaviour: the desire to conform with peers who prefer not to test increases the number of households in which no one tests by around one third. At the other extreme, the proportion of households in which everyone tests is 6 percentage points higher in the actual household data than in the placebo data, corresponding to an increase of around three quarters. This increase likely reflects the combined effects of externalities and social norms on students who would not have tested if left to their own devices.

5.3 Total testing

By focusing on the extremes of the distribution, we have established that students' testing decisions are influenced by social norms.

Moving away from the extremes, the share of households in which more than half of the students test is significantly higher in the actual households than in the placebo house-

Table 6: Comparison of proportions at the two extremes

	All households		Forced singleton placebo		Freshers & forced singletons	
Number testing:	All (1)	No-one (2)	All (3)	No-one (4)	All (5)	No-one (6)
Actual household	0.14 (0.00)	0.30 (0.01)	0.14 (0.00)	0.30 (0.01)	0.11 (0.01)	0.37 (0.01)
Placebo household	0.07 (0.00)	0.20 (0.00)	0.08 (0.00)	0.20 (0.00)	0.07 (0.00)	0.24 (0.01)
N	46,406	46,406	50,964	50,964	13,494	13,494

Notes: Coefficients of linear regressions. Dependent variables are dummy variables equal to 1 if all test in the household (columns 1, 3 and 5) and if no-one tests in the household (columns 2, 4 and 6). Independent variables are dummies for whether the household is actual or placebo (no constant is included). Unit of observation is household-week. Standard errors are adjusted for serial (overtime) correlation within households. Sample (i) Time: Lent and Easter term 2021 in all columns (ii) Actual households: all households with 2-9 members (columns 1-4); fresher households with 2-9 members (columns 5-6); (iii) Placebo households constructed from students in single accommodation (columns 1-2), and students in multiperson households who were found themselves alone (columns 3-6).

holds.[34] As shown in the model developed in Section 3, this pattern is consistent with both externalities and social norms.

Given that social norms can operate in both directions, it is natural to ask whether grouping students increases or decreases total testing. Although we do not observe an experimental counterfactual, we can compare testing behaviour among students who live with others to that of students who live alone. The results are reported in Table 7.

Following the approach in Section 5.2, we first define students living alone as those in single-student households, regardless of the reason (columns 1 and 2). Under this definition, we find no significant effect of living in a group on testing. Second, to mitigate concerns about self-selection into living alone, we redefine living alone as being left on one's own in a multi-person household because housemates did not return to Cambridge after the Christmas break in early 2021 (columns 3 and 4). Under this definition, we find a significant negative effect of living alone on the likelihood of testing, of around 5–6 percentage points.

At first glance, this effect may appear small, especially given that the peer-effect regressions show that a student's probability of testing increases by around 70 percentage points when

34. This is illustrated in Figure 5 and holds across all household sizes; see the figures in Appendix **??**.

all household members begin testing. The reason is that, as demonstrated in Section 5.2, household effects—like peer effects—operate in both directions due to heterogeneity in social norms. In some households, students who would not have tested on their own do so to conform with the prevailing norm, while in others, students who would otherwise have tested abstain in order to conform.

The results in Table 7 underscore the importance of distinguishing social norms from externalities. Since total testing is higher when students live with others than when they live alone, one might be tempted to conclude that grouping individuals is an effective way to increase prosocial behaviour such as testing. However, this aggregate effect masks substantial heterogeneity, with different households developing opposing social norms. As a result, whether aggregate testing increases or decreases is a priori ambiguous, as we show in Section 3.4. This stands in sharp contrast to a setting with only positive externalities, in which grouping students would unambiguously increase aggregate testing.

Table 7: Effect of not being in a household on total testing

	Single accommodation		Forced singleton	
	All (1)	Freshers (2)	All (3)	Freshers (4)
Living alone	-0.016	0.016	-0.047***	-0.059**
	(0.015)	(0.026)	(0.013)	(0.025)
Controls	✓	✓	✓	✓
Week effect	✓	✓	✓	✓
N	238,243	53,458	145,574	34,132

Dependent variable: dummy variable =1 if the student tests. Unit of observation: student-week. Estimator is OLS. Standard errors are in parentheses and are clustered at the household level. The regressions include week fixed effects and controls. Independent variable: Columns 1-2: a dummy variable =1 if the student lives in single accommodation; Columns 3-4 a dummy variable =1 if the student lives alone because his housemates did not return to Cambridge. Sample: columns 1 - all students; columns 2 - freshers; column 3 - all students excluding those in single accommodation; column 4 - freshers excluding those in single accommodation. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

6 Conclusions

We develop a theoretically-motivated empirical test that can separate social norms from externalities in observational data that display peer effects, overcoming a well-known non-identifiability result. We do so in the important context of prosocial choices, and, more specifically, student participation in a large-scale asymptomatic testing programme operated by the University of Cambridge during the Covid-19 pandemic. We show that within student households, testing decisions are subject to strong positive peer effects, implying strategic complementarity. This is consistent with both externalities and social norms. To distinguish between these two mechanisms, we implement a novel theory-motivated test based on the analysis of the distribution of testing outcomes across households. By comparing the observed distribution to what we would expect if students actions were independent, we show that the observed peer effects in student behaviour cannot be explained by externalities alone and that social norms must have been present.

We show that different households developed different social norms: in some households, the social norm is to test, whilst in other households it is not to test. This is consistent with the theoretical literature that argues that social norms are often subject to multiple equilibria. The presence of diverging social norms cannot be picked up by standard econometric tools, but has important policy implications.

References

Aggarwal, Dinesh, Ben Warne, et al. 2022. "Genomic epidemiology of SARS-CoV-2 in a UK university identifies dynamics of transmission." *Nature Communications* 31 (751).

Bénabou, Roland, and Jean Tirole. 2006. "Incentives and Prosocial Behavior." *American Economic Review* 96 (5): 1652–1677.

Bernheim, B. Douglas. 1994. "A Theory of Conformity." *Journal of Political Economy* 102 (5): 841–877.

Bhattacharya, Debopam, Pascaline Dupas, and Shin Kanaya. 2024. "Demand and Welfare Analysis in Discrete Choice Models with Social Interactions." Advance access published 11 May 2023, *The Review of Economic Studies* 91, no. 2 (March): 748–784. https://doi.org/10.1093/restud/rdad053. https://doi.org/10.1093/restud/rdad053.

Bicchieri, Cristina, and Alex Chavez. 2010. "Behaving as Expected: Public Information and Fairness Norms." *Journal of Behavioral Decision Making* 23 (2): 161–178. https://doi.org/10.1002/bdm.648.

Bicchieri, Cristina, Eugen Dimant, Simon Gächter, and Daniele Nosenzo. 2022. "Social proximity and the erosion of norm compliance." *Games and Economic Behaviour* 132:59–72. https://doi.org/10.1016/j.geb.2021.11.012.

Bicchieri, Cristina, Eugen Dimant, Michele Gelfand, and Silvia Sonderegger. 2023. "Social norms and behavior change: The interdisciplinary research frontier." *Journal of Economic Behavior & Organization* 205:A4–A7. https://doi.org/10.1016/j.jebo.2022.11.007.

Bicchieri, Cristina, Ryan Muldoon, and Alessandro Sontuoso. 2023. "Social Norms." In *The Stanford Encyclopedia of Philosophy,* Winter 2023, edited by Edward N. Zalta and Uri Nodelman. Stanford, CA: Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/entries/social-norms/.

Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Rajshri Jayaraman. 2015. "Linear Social Interactions Models." *Journal of Political Economy* 123 (2): 444–496. https://doi.org/10.1086/679496.

Boucher, Vincent, Yann Bramoullé, Habiba Djebbari, and Bernard Fortin. 2014. "Do Peers Affect Student Achievement? Evidence from Canada Using Group Size Variation." *Journal of Applied Econometrics* 29 (1): 91–109. https://doi.org/10.1002/jae.2306.

Boucher, Vincent, and Bernard Fortin. 2016. "Some Challenges in the Empirics of the Effects of Networks." In *The Oxford Handbook of the Economics of Networks,* 277–302. Oxford, UK: Oxford Handbooks.

Boucher, Vincent, Michelle Rendall, Philip Ushchev, and Yves Zenou. 2024. "Toward a General Theory of Peer Effects." *Econometrica* 92 (2): 543–565. https://doi.org/10.3982/ECTA21048.

Brock, William A., and Steven N. Durlauf. 2001. "Discrete Choice with Social Interactions." *Review of Economic Studies* 68:235–260.

———. 2007. "Identification of Binary Choice Models with Social Interactions." *Journal of Econometrics* 140:52–75.

Bursztyn, Leonardo, Georgy Egorov, and Robert Jensen. 2019. "Cool to be Smart or Smart to be Cool? Understanding Peer Pressure in Education." *The Review of Economic Studies* 86 (4): 1487–1526.

Bursztyn, Leonardo, and Robert Jensen. 2017. "Social Image and Economic Behavior in the Field: Identifying, Understanding, and Shaping Social Pressure." *Annual Review of Economics* 9:131–153.

DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature* 47 (2): 315–372.

DellaVigna, Stefano, John A. List, and Ulrike Malmendier. 2012. "Testing for Altruism and Social Pressure in Charitable Giving." *Quarterly Journal of Economics* 127 (1): 1–56. https://doi.org/10.1093/qje/qjr050.

DellaVigna, Stefano, John A. List, Ulrike Malmendier, and Gautam Rao. 2017. "Voting to Tell Others." *Review of Economic Studies* 84 (1): 143–181. https://doi.org/10.1093/restud/rdw056.

Elster, Jon. 1989. "Social Norms and Economic Theory." *Journal of Economic Perspectives* 3 (4): 99–117. https://doi.org/10.1257/jep.3.4.99.

Fang, Ximeng, Timo Freyer, Chui-Yee Ho, Zihua Chen, and Lorenz Goette. 2022. "Prosociality predicts individual behavior and collective outcomes in the COVID-19 pandemic." *Social Science & Medicine* 308:115192. https://doi.org/10.1016/j.socscimed.2022.115192. https://doi.org/10.1016/j.socscimed.2022.115192.

Fehr, Ernst, and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114 (3): 817–868. https://doi.org/10.1162/003355399556151.

Funk, Patricia. 2010. "Social Incentives and Voter Turnout: Evidence from the Swiss Mail Ballot System." *Journal of the European Economic Association* 8 (5): 1077–1103.

Gächter, Simon, Daniel Nosenzo, and Martin Sefton. 2013. "Peer Effects in Pro-Social Behavior: Social Norms or Social Preferences." *Journal of the European Economic Association* 11 (3): 548–573.

Goette, Lorenz, and Egon Tripodi. 2021. "Social Influence in Prosocial Behavior: Evidence from a Large-scale Experiment." *Journal of the European Economic Association* 19 (4): 2373–2398.

Granovetter, Mark. 1978. “Threshold Models of Collective Behavior.” *American Journal of Sociology* 83 (6): 1420–1443.

Krupka, Erin L., and Roberto A. Weber. 2013. “Identifying Social Norms Using Coordination Games: Why Does Dictator Game Sharing Vary?” *Journal of the European Economic Association* 11 (3): 495–524. https://doi.org/10.1111/jeea.12006.

Levitt, Steven D. 2020. *I’m Not as Childlike as I’d Like to Be.* People I (Mostly) Admire, Extra Episode, 9 October 2020, Freakonomics Radio Podcast Episode transcript. Accessed September 25, 2024. https://freakonomics.com/podcast/steve-levitt-im-not-as-childlike-as-id-like-to-be-bonus/.

Manski, Charles. 1993. “Identification of Endogenous Social Effects: The Reflection Problem.” *The Review of Economic Studies* 60 (3): 531–542.

Mas, Alexandre, and Enrico Moretti. 2009. “Peers at Work.” *American Economic Review* 99 (1): 112–145.

Matheson, Nicholas, Ben Warne, Michael P. Weekes, and Patrick H. Maxwell. 2021. “Mass testing of university students for covid-19.” *British Medical Journal* 375 (2388).

Postlewaite, Andrew. 2011. “Social Norms and Social Assets.” *Annual Review of Economics* 3:239–259. https://doi.org/10.1146/annurev-economics-061109-080158.

Smith, Sarah, Frank Windmeijer, and Edmund Wright. 2015. “Peer Effects in Charitable Giving: Evidence from the (Running) Field.” *The Economic Journal* 125:1053–1071.

University of Cambridge. 2020. *Stay Safe, Cambridge University Asymptomatic testing – further information.*

———. 2022. *How mass testing helped limit the spread of COVID-19 at the University of Cambridge.* Accessed September 19, 2024. https://www.cam.ac.uk/stories/screeningprogramme.

Young, H. Peyton. 2015. “The Evolution of Social Norms.” *Annual Review of Economics* 7:359–387. https://doi.org/10.1146/annurev-economics-080614-115322.

A Appendix: Testing instructions

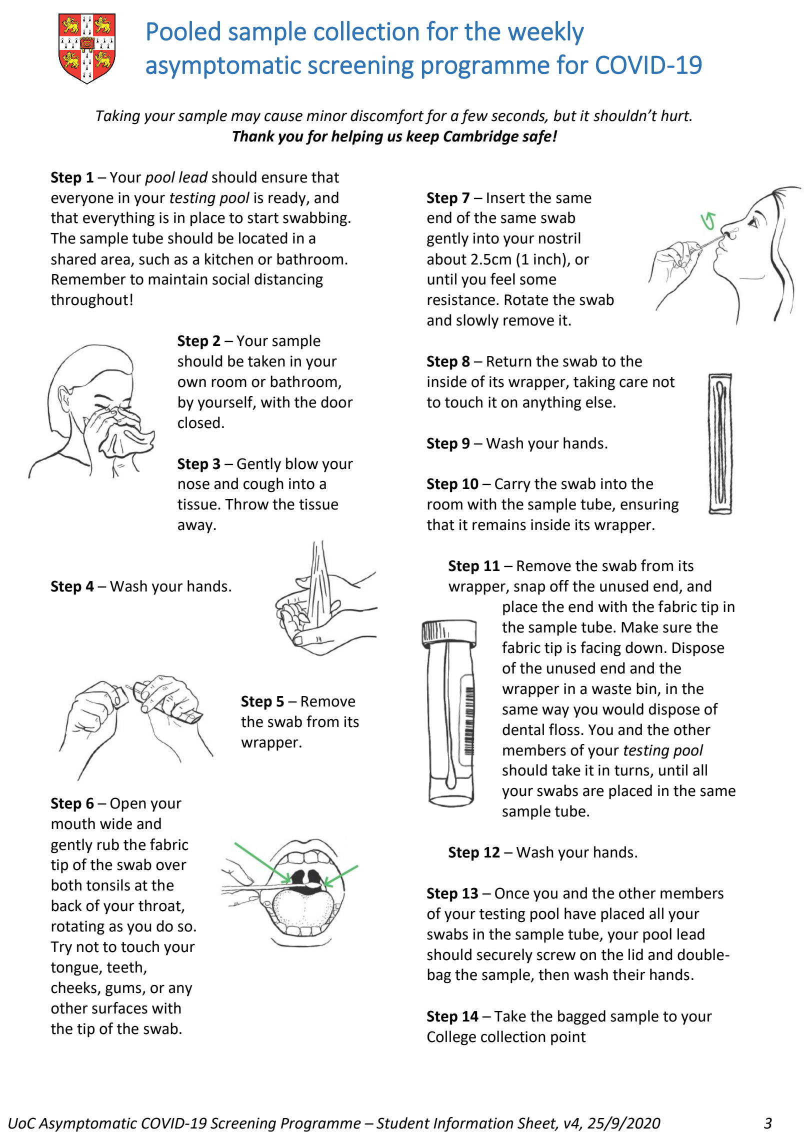

Pooled sample collection for the weekly asymptomatic screening programme for COVID-19

Taking your sample may cause minor discomfort for a few seconds, but it shouldn't hurt.
Thank you for helping us keep Cambridge safe!

Step 1 – Your *pool lead* should ensure that everyone in your *testing pool* is ready, and that everything is in place to start swabbing. The sample tube should be located in a shared area, such as a kitchen or bathroom. Remember to maintain social distancing throughout!

Step 2 – Your sample should be taken in your own room or bathroom, by yourself, with the door closed.

Step 3 – Gently blow your nose and cough into a tissue. Throw the tissue away.

Step 4 – Wash your hands.

Step 5 – Remove the swab from its wrapper.

Step 6 – Open your mouth wide and gently rub the fabric tip of the swab over both tonsils at the back of your throat, rotating as you do so. Try not to touch your tongue, teeth, cheeks, gums, or any other surfaces with the tip of the swab.

Step 7 – Insert the same end of the same swab gently into your nostril about 2.5cm (1 inch), or until you feel some resistance. Rotate the swab and slowly remove it.

Step 8 – Return the swab to the inside of its wrapper, taking care not to touch it on anything else.

Step 9 – Wash your hands.

Step 10 – Carry the swab into the room with the sample tube, ensuring that it remains inside its wrapper.

Step 11 – Remove the swab from its wrapper, snap off the unused end, and place the end with the fabric tip in the sample tube. Make sure the fabric tip is facing down. Dispose of the unused end and the wrapper in a waste bin, in the same way you would dispose of dental floss. You and the other members of your *testing pool* should take it in turns, until all your swabs are placed in the same sample tube.

Step 12 – Wash your hands.

Step 13 – Once you and the other members of your testing pool have placed all your swabs in the sample tube, your pool lead should securely screw on the lid and double-bag the sample, then wash their hands.

Step 14 – Take the bagged sample to your College collection point

UoC Asymptomatic COVID-19 Screening Programme – Student Information Sheet, v4, 25/9/2020 3

Figure 7: Testing instructions for students

B Appendix: Distributions by term and household size

In section 5.2 we used the visual example of six student households in Lent and Easter terms to demonstrate the presence of social norms. The social norms are identified because the distribution of testing students across households has both a higher share of households where all test, and a higher share of households where no-one tests compard to the independence counterfactual. Figures 8 and 9 show similar pattern across most household sizes in Lent and Easter 2020-21 and Michaelmas 2021-22, respectively.

Figure 8: Number testing. All students, Lent and Easter 2021

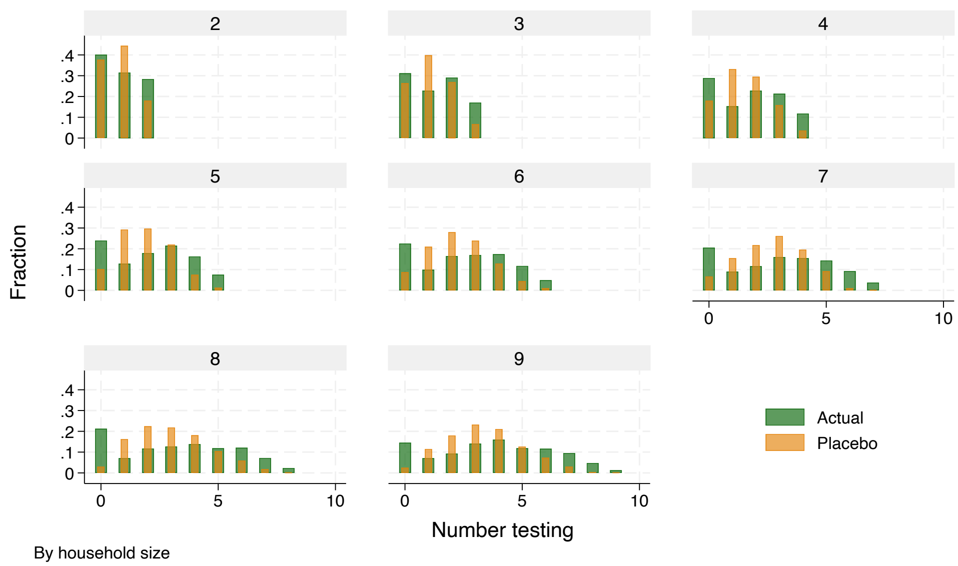

Figure 9: Number testing. All students, Michaelmas 2021

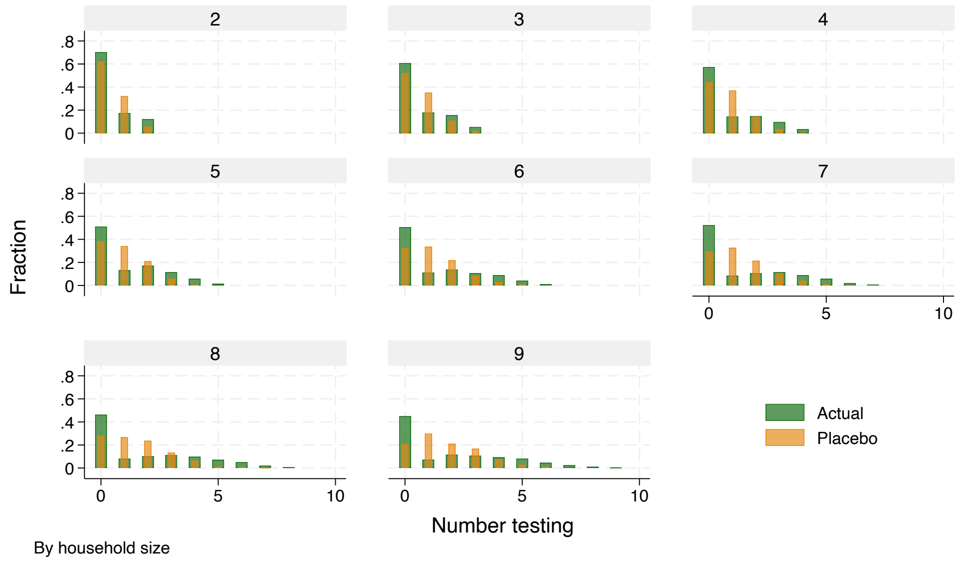

C Appendix: Negative externalities

Although we do not think that significant negative externalities are plausible in our context, and their presence is not in line with our empirical results, we extend our model to them and show that validity of our test is unchanged.

We take a concrete example: As infection rates are correlated, someone might feel that the social benefit of them testing is lower if others are already testing. We model this as a decline in i's benefit from testing when a higher proportion of the household tests ($\bar{t}_h$). Now, i's benefit from testing is $b_i - \tilde{d}(\bar{t}_h)$, where $\tilde{d}(\bar{t}_h)$ is the externality, such that $\frac{\partial \tilde{d}}{\partial \bar{t}_h} > 0$, $\tilde{d}(0) = 0$ and $\underline{b} \leq \tilde{d}(1) > 0$.

We introduce this modification into our model with externalities ($k = E$) in section 3.3. Now, testing decisions are strategic complements when $\frac{\partial \tilde{d}}{\partial \bar{t}_h} + \frac{\partial \tilde{c}}{\partial \bar{t}_h} < 0$, i.e. cost externalities dominate benefit externalities, and strategic substitutes otherwise.

The Nash equilibrium household-level outcomes are summarized in Lemma 7.

Lemma 7 (Negative externalities) *The Nash equilibrium household-level outcome for household h of the simultaneous-move game with positive externalities can be characterised as follows:*

- *All students in H_h test ($j = A$): $n_h^{**} = N_h$ if and only if $\underline{b} \geq \tilde{d}(1) - \tilde{c}(1)$.*
- *No student in H_h tests ($j = N$): $n_h^{**} = 0$ if and only if $\bar{b}_h < c$.*
- *Some students in H_h test ($j = M$): $n_h^{**} \in \{1, \ldots, N_h - 1\}$ if and only if $\bar{b}_h \geq \tilde{d}(1) + \tilde{c}(1)$ and $\underline{b}_h < c$.*

Intuitively, cost externalities can persuade some of those who did not test under independence to test now because their costs are lower when others are testing. Conversely, benefit externalities can persuade some of those who tested under independence not to test now *as long as* there are others who are testing in their household. However, if no-one else in the household is testing, there are no externalities, and so an individual's testing decisions are the same as under independence. Hence the condition for no-one tests to be a Nash are the same as under independence.

The share of households with the three outcomes will now depend on the relative strength of the two externalities and which of the multiple equilibria are played. However, we can still put some bounds on the share of households where no-one tests relative to the independence case:

Proposition 8 *The distribution of equilibrium outcomes with both positive and negative externalities ($k = E$) compared to the distribution with independent payoffs ($k = I$) must satisfy*

$$S_N^E \leq S_N^I$$

The proposition holds since the condition for the 'no-one tests' equilibrium is the same as under the independence case, however, one of the other types of equilibria (all or some test) will (weakly) expand (which one depends on which externality dominates). This means that the share households where no-one tests cannot be bigger than under independence, and may be smaller due to other equilibria.

Hence, the result of Proposition 6, and our test for social norms remains unchanged: if the observed empirical distribution involves $\hat{S}_N > \hat{S}_N^I$, this implies that the data-generating process must have involved social norms.

D Appendix: Social norms in Michaelmas term

Table 8 shows the test for presence of social norms in Michaelmas term 2021, the last term of the Programme. The results are qualitatively the same as we found in Lent and Easter terms (section 5.2). Quantitatively, participation in the Programme declined over time, and so the likelihood of no-one testing is higher, whilst that of everyone testing is lower than in previous terms. This is true for both decisions under independence and actual decisions.

Table 8: Probability of 'No-one tests' and 'All test', Michaelmas term 2021

	All households		Fresher households	
	All test	No-one tests	All test	No-one tests
	(1)	(2)	(3)	(4)
Actual household	0.10	0.77	0.16	0.69
	(0.01)	(0.01)	(0.02)	(0.02)
Placebo household	0.07	0.64	0.11	0.58
	(0.01)	(0.01)	(0.01)	(0.02)
Observations	34731	34731	9585	9585

Notes: Coefficients of linear regressions. Dependent variables are dummy variables equal to 1 if all test in the household (columns 1 and 3) and if no-one test in the household (columns 2 and 4). Independent variables are dummies for whether the household is actual or placebo (no constant is included). Unit of observation is household-week. Standard errors are adjusted for serial (overtime) correlation within households. Sample (i) Time: Michaelmas term 2021; (ii) Actual households: all households with 2-9 members (columns 1-2); fresher households with 2-9 members (columns 3-4); (iii) Placebo households constructed from students in single accommodation.