



UNIVERSITÀ
DEGLI STUDI
FIRENZE

FLORE

Repository istituzionale dell'Università degli Studi di Firenze

Exploiting network information to disentangle spillover effects in a field experiment on teens' museum attendance

Questa è la versione Preprint (Submitted version) della seguente pubblicazione:

Original Citation:

Exploiting network information to disentangle spillover effects in a field experiment on teens' museum attendance / Silvia Noirjean, Marco Mariani, Alessandra Mattei, Fabrizia Mealli. - ELETTRONICO. - (2020), pp. 0-0. [10.48550/arXiv.2011.11023]

Availability:

The webpage <https://hdl.handle.net/2158/1340671> of the repository was last updated on 2023-10-30T15:50:42Z

Published version:

DOI: 10.48550/arXiv.2011.11023

Terms of use:

Open Access

La pubblicazione è resa disponibile sotto le norme e i termini della licenza di deposito, secondo quanto stabilito dalla Policy per l'accesso aperto dell'Università degli Studi di Firenze (<https://www.sba.unifi.it/upload/policy-oa-2016-1.pdf>)

Publisher copyright claim:

La data sopra indicata si riferisce all'ultimo aggiornamento della scheda del Repository FloRe - The above-mentioned date refers to the last update of the record in the Institutional Repository FloRe

(Article begins on next page)

Exploiting Network Information to Disentangle Spillover Effects in a Field Experiment on Teens' Museum Attendance

Silvia Noirjean¹, Marco Mariani², Alessandra Mattei^{1,3}, and
Fabrizia Mealli^{1,3}

¹*Department of Statistics, Computer Science, Applications "G. Parenti", University of Florence*

²*IRPET - Regional Institute for Economic Planning of Tuscany*

³*Florence Center for Data Science*

A key element in the education of youths is their sensitization to historical and artistic heritage. We analyze a field experiment conducted in Florence (Italy) to assess how appropriate incentives assigned to high-school classes may induce teens to visit museums in their free time. Non-compliance and spillover effects make the impact evaluation of this clustered encouragement design challenging. We propose to blend principal stratification and causal mediation, by defining sub-populations of units according to their compliance behavior and using the information on their friendship networks as mediator. We formally define principal natural direct and indirect effects and principal controlled direct and spillover effects, and use them to disentangle spillovers from other causal channels. We adopt a Bayesian approach for inference.

Keywords: *Bayesian inference; causal inference with network interference; clustered encouragement design; mediation analysis; principal stratification*

1 Introduction

Over the last years, there has been an increasing interest in causal studies where units are part of networks (e.g., groups of friends). In this type of studies, social or physical interactions among units are a major concern, and thus, the no-interference assumption, under which the outcome of each unit cannot depend on the treatment assigned to others (Cox, 1958), is no longer plausible.

The literature on causal inference in the presence of networks has developed considerably in recent years (e.g. Goldsmith-Pinkham & Imbens, 2013; VanderWeele & An, 2013; Kang & Imbens, 2016; Ogburn et al., 2017; Forastiere et al., 2018; Forastiere, Airoidi, & Mealli, 2021). The intuition characterizing this literature is that interference between units can be conveyed by the specific social links that tie these units together into a network, irrespective of treatment assignment. With this in mind, it is as if each unit was exposed to multiple treatments: its own treatment and the treatments of the other units in its network. The latter may give rise to spillover effects (Del Prete et al., 2019).

Interference naturally arises in clustered encouragement designs (CEDs), where encouragement is randomized at the level of a cluster (e.g., class) of subjects (e.g., students) but compliance with the assigned encouragement is often at the individual level. CEDs are widely adopted, especially when randomization of the receipt of treatment is not feasible for ethical or practical reasons and the assignment of encouragement at the level of clusters appears preferable for administrative convenience, cost reduction, ethical issues and for improving treatment adherence. The literature abounds of studies designed as CEDs with individual noncompliance (e.g., Frangakis et al., 2002; Morris et al., 2004; Forastiere et al., 2016).

Clusters define a special network, where sets of units belonging to the same cluster are likely to share neighbors, but have no links with units belonging to different clusters. Therefore, in CEDs social interactions naturally occur among units belonging to the same group, and we can reasonably expect that among subjects who interact with one another, the treatment actually received by one subject may affect other subjects' outcomes, gener-

ating interference or spillover effects (Sobel, 2006; Hong & Raudenbush, 2006; Hudgens & Halloran, 2008). The presence of interference gives rise to intricate causal mechanisms, and disentangling them may provide precious information for researchers and policymakers to understand how an intervention works and how it can be improved.

A CED with individual noncompliance is the study design underlying our motivating study on teens’ museum attendance. It is a field experiment based on an encouragement design clustered at the level of school classes, conducted in Florence (Italy) and designed to assess how appropriate incentives may prompt teens to visit museums in their free time. Throughout the text we refer to this study as the Florentine study. A unique feature of this study is the availability of information about the friendship ties between the teens belonging to the same class, which makes the study a *CED with a hierarchical network structure*, with a high-level network defined by the clustered structure and a low-level network defined by the friendship structures within each cluster. The study was previously analyzed by Lattarulo et al. (2017) and Forastiere, Lattarulo, et al. (2021), but they neglect information on friendship ties, and thus, provide only a partial picture of the causal mechanisms by which the designed incentives affect teens’ museums attendance.

We make important substantive contributions by adopting an innovative approach to the analysis of CEDs with individual noncompliance under network interference. Teens’ museums attendance is a relevant topic because, while teens are likely to visit museums during school or family trips, museums are seldom their choice when it is up to them to choose where to spend their free time (e.g. Hughes & Moscardo, 2019). In addition to granting free or low-price entrance to youths (Cellini & Cuccia, 2018), many museums are attempting to portray an image of educational and entertaining institutions and devising strategies to attract teens and young adults towards a tailored visit experience (Lattarulo et al., 2017; Manna & Palumbo, 2018). The hope is that such experience can raise further engagement with museums and cultural heritage in general (Kisida et al., 2014). In the Florentine experiment, classes of students are randomized to three forms of encouragement

aimed to prompt students to visit Palazzo Vecchio, one of the most famous museums in the city. Students belonging to a first group of classes receive a flyer with basic information about Palazzo Vecchio, including its opening hours. Students in a second group of classes receive both the flyer and a short presentation conducted by an art expert. The goal of the presentation is to enhance students’ curiosity about museum visits, in general, and Palazzo Vecchio, in particular. Students in a third group of classes, in addition to the flyer and the presentation, receive a non-financial reward in the form of extra-credit points towards their school grade. The treatment is the actual visit to Palazzo Vecchio, which can be performed in the two months after the encouragement. The outcome of interest is the number of museum visits performed in the follow-up period, that is, between two and eight months from the encouragement.

The stronger the encouragement, the higher the propensity of a student to visit Palazzo Vecchio is expected to be. However, students’ compliance with the cluster-level encouragement is imperfect, meaning that there are both students in classes assigned to the flyer who visit Palazzo Vecchio as well as students in classes assigned to the reward who do not visit Palazzo Vecchio. We deal with the presence of noncompliance using the principal stratification (PS) framework (Frangakis & Rubin, 2002): we define stratum-specific causal effects, named principal causal effects, which are effects for specific latent subpopulations, defined by the joint potential compliance statuses under the three encouragement conditions. The role of PS in encouragement designs to draw inference on intention-to-treat effects within principal strata is uncontroversial (e.g., Angrist et al., 1996; Imbens & Rubin, 1997; Hirano et al., 2000; Frangakis et al., 2002).

How to deal with network interference in CEDs is still an open research issue. In principle, a number of causal pathways may take place within CEDs. The first path is the effect of the encouragement on a unit’s outcome passing through that unit’s treatment uptake. The second path is through spillover effects, i.e., changes in a unit’s outcome originating from the treatment uptake by other units, which naturally arise among friends. The third path

is the direct effect of the encouragement passing neither through the individual treatment uptake nor through spillovers, which ends up pooling further causal channels.

Two recent studies by Forastiere et al. (2016) and Forastiere, Lattarulo, et al. (2021) provide insightful contributions on the causal mechanisms in action in CEDs. Forastiere et al. (2016) propose to use the principal stratification framework to disentangle distinct causal mechanisms in CEDs, viewing the treatment variable as a mediating variable and defining causal effects conditional on principal stratum membership based on hypothetical interventions on the treatment uptake. Forastiere, Lattarulo, et al. (2021) show how intention-to-treat effects within principal strata can be interpreted to get valuable information on different causal mechanisms for different types of units. Both Forastiere et al. (2016) and Forastiere, Lattarulo, et al. (2021) find evidence that encouragement and treatment effects in their CEDs are blended with spillovers, and highlight that disentangling these effects may be uneasy. We contribute to this literature providing an innovative method for untying encouragement, treatment and spillovers in CEDs with a hierarchical network structure, which capitalize on the distinguish structure of our data: the availability of within cluster network data. Specifically, our dataset includes information on the friendship network within the class of each student participating in the study. We use this information to construct a new variable defined as the proportion of friends who actually take the treatment, that is, visit Palazzo Vecchio. Our key insight is to use this variable as a mediator and conduct principal stratum mediation analysis to disentangle distinct causal mechanisms. We define principal stratum natural direct and indirect effects and principal stratum controlled direct effects, which are natural direct and indirect effects and controlled direct effects (Robins & Greenland, 1992; Pearl, 2001) for subpopulations of units defined by the principal stratification with respect to the individual receipt of treatment. We also introduce new causal estimands, which we refer to as principal controlled spillover effects, defined as comparisons within principal strata of hypothetical potential outcomes for the response variable under a fixed level of the encouragement and different values of the mediator. We interpret principal

natural indirect effects and principal controlled spillover effects as spillover effects, which may be heterogeneous across different types of units defined by the compliance statuses. Principal natural and controlled direct effects, instead, can be viewed as pure effect of the encouragement or as a blend of the encouragement and the treatment/experience effects, depending on the principal stratum and the encouragement conditions being compared.

Inference on the causal estimands of interest is conducted under latent ignorability assumptions of the mediator using a Bayesian approach. We introduce two latent ignorability assumptions of the mediator: a latent ignorability assumption of the observed mediator, implying that the observed proportion of friends visiting Palazzo Vecchio are independent of potential outcomes for the number of museum visits conditional on principal stratum membership given the observed value of the encouragement and individual level and network level pretreatment covariates; and a latent ignorability assumption of the potential mediator, implying that potential outcomes for the proportion of friends visiting Palazzo Vecchio are independent of potential outcomes for the number of museum visits conditional on principal stratum membership given observed individual level and network level pretreatment covariates. Under latent ignorability assumptions, ignorability of the mediator only holds after conditioning on principal stratum membership, which is a latent variable, only partially observed for each unit. The latent nature of the principal strata complicates inference. We deal with inferential issues using a Bayesian approach, which does not require full identification (e.g. Gustafson, 2010). We specify flexible parametric models for the principal stratum membership given individual- and network-level covariates and for the potential outcomes conditional on principal stratum membership and individual- and network-level covariates.

Our work provide a cutting-edge innovative approach to the analysis of CEDs, which smartly blend concepts and tools from principal stratification, mediation analysis and causal inference with network data.

The article is organized as follows. Section 2 introduces the experimental design and the data. Section 3 describes the methodology. Section 4 reports the results of the analysis and

Section 5 concludes.

2 Experimental Design and Data

The clustered encouragement field experiment we re-analyze was run in Florence, Italy, at three different points in time during 2014. It involved $J = 15$ classes, indexed by $j = 1, \dots, 15$, from 3 different high schools in the city. Let N_j be the number of students in class j , $j = 1, \dots, 15$, and let $N = \sum_{j=1}^J N_j$ be the total number of students participating in the study: $N = 266$. Finally, let C_{ij} be an indicator for student i 's class membership: $C_{ij} = 1$ if student i is in class j and zero otherwise, $i = 1, \dots, N_j$ and $j = 1, \dots, J$. All high schools were of the same type and all students in the classes participating in the study were aged 17-18. Each of the 15 classes was randomly assigned to one of three incremental encouragements, in groups of five. All students, irrespective of the encouragement their class was assigned to, were offered the opportunity to visit individually the Florentine Museum of Palazzo Vecchio outside of school hours.

Under the weakest encouragement, which we name *flyer*, students received a flyer containing basic information about Palazzo Vecchio (opening hours and a brief description of the museum), and a short text written by the experimenters, briefly stating the importance of museum attendance. Under the intermediate encouragement, named *presentation*, in addition to the flyer and text, students received a short presentation conducted by an art expert from the museum: the aim of the presentation was to enhance students' curiosity about museum visits, in general, and Palazzo Vecchio, in particular, and to portray museum attendance as an intriguing and entertaining experience. Under the strongest encouragement, which we name *reward*, in addition to the flyer, text and presentation, students were promised a non-financial reward in the form of extra-credit points towards their final school grade if they actually visited Palazzo Vecchio within two months after assignment. Let Z_j denote the encouragement which class j is assigned to: $Z_j=1$ if the class is assigned to the

flyer, $Z_j=2$ if it is assigned to the *presentation* and $Z_j=3$ if it is assigned to the *reward*.

The field experiment was implemented through three visits to the classes, whose timing and contents are described in what follows. At time $t = 1$, all students were surveyed about their background characteristics and interviewed about their own friendship network within the class. In particular, they were asked to mention all the classmates they consider their friends. Then, classes were randomly assigned to the three encouragement levels described above. All students were offered a free visit to Palazzo Vecchio to be made within two months.

At time $t = 2$, two months after the assignment, information about whether students had visited or not Palazzo Vecchio was collected. Let M_{ij} be a binary variable, equal to 1 if student i in class j visits Palazzo Vecchio in the two months after the assignment, and zero otherwise. The variable M_{ij} is the treatment of interest.

A follow-up period of 6 months, starting from time $t = 2$ was considered. At the end of the study, that is, at time $t = 3$ (6 months after $t = 2$), information on the number of visits each student made to other museums in the follow-up period was collected. Let Y_{ij} be a count variable indicating the number of museum visits made by student i in class j during the follow-up period: Y_{ij} is the outcome of primary interest. See Lattarulo et al. (2017) and Forastiere, Lattarulo, et al. (2021) for a more detailed description of the three forms of encouragement and a comprehensive analysis of the substantive results of the field experiment.

Lattarulo et al. (2017) use classes as units of analysis and resort to randomization inference techniques to conduct an intention-to-treat analysis of the differential effects of the three encouragements on classroom-level museum attendance. Forastiere, Lattarulo, et al. (2021), instead, use students as units of analysis considering the students' choice to visit Palazzo Vecchio as an endogenous treatment. Under the partial interference assumption (e.g., Sobel, 2006), they use PS to investigate the heterogeneity of causal effects across different latent subgroups of units. Although they find that, for all these subgroups, encouragement and/or

treatment effects are mixed with spillovers, disentangling these effects is beyond the aim of their study. Our methodological contribution fills this gap using information on individual friendship networks, which is not exploited in Forastiere, Lattarulo, et al. (2021).

Let us denote students' friendship network as a pair $(\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ is the set of nodes (subjects) and $\mathcal{E} = \{E_{ik}, i, k = 1, \dots, N_j; j = 1, \dots, J\}$ is the set of edges (friendship ties) that links the subjects; in particular, $E_{ik} = E_{ki} = 1$ if subject i declared subject k as a friend and/or viceversa, and $E_{ik} = E_{ki} = 0$ otherwise. It is worth noting that, by design, $E_{ik} = E_{ki} = 0$ for each pair of students i and k belonging to two different classes. Therefore we are considering a binary and undirected network: "binary" because we assign an equal weight to all the friendships, and "undirected" because an influence between subject i and subject k is plausible if at least one of them declared the other as a friend. For each student i in class j , let $\mathcal{N}_{ij}$ denote the neighborhood of student i , comprising student i 's friends in class j .

For each student i in class j , $i = 1 \dots, N_j$, $j = 1, \dots, J$, we define a new variable $S_{\mathcal{N}_{ij}}$ as the proportion of student i 's friends who visit Palazzo Vecchio. $S_{\mathcal{N}_{ij}}$ is a summary of $M_{i'j}$ for all $i' \neq i \in \mathcal{N}_{ij}$:

$$S_{\mathcal{N}_{ij}} = \frac{1}{|\mathcal{N}_{ij}|} \sum_{i' \neq i \in \mathcal{N}_{ij}} M_{i'j}$$

where $|\mathcal{N}_{ij}|$ is the number of student i 's friends, that is, the degree of student i .

For each student i in class j , $i = 1 \dots, N_j$, $j = 1, \dots, J$, we observe a vector of K covariates, $\mathbf{X}_{ij}$. This vector is composed by two sub-vectors: $\mathbf{X}_{ind_{ij}}$ and $\mathbf{X}_{\mathcal{N}_{ij}}$, which include individual and neighborhood-level covariates, respectively. The sub-vector $\mathbf{X}_{\mathcal{N}_{ij}}$ includes different types of neighborhood-level covariates: some of them are a synthesis of individual-level covariates in the neighborhood, such as, e.g., the proportion of male friends or the average number of museum visits made by friends prior to the experiment; others represent the structure of the neighborhood, such as the degree (number of ties) of a unit or the average degree of a unit's friends.

2.1 Descriptive statistics

Table 1 reports some descriptive statistics for the sample of the 266 students, their friends and the structure of their individual friendship network, grouped by class assignment, Z_j . As it is often the case with small-scale field experiments, background characteristics are not

Table 1

The museum study: descriptive statistics of the individual and friends' level background covariates (proportions or means)

	<i>Flyer</i>	<i>Presentation</i>	<i>Reward</i>	<i>Overall</i>
<i>INDIVIDUAL BACKGROUND CHARACTERISTICS</i>				
Male (1/0)	0.19	0.29	0.53	0.34
Already visited Palazzo Vecchio (1/0)	0.66	0.75	0.72	0.71
No. of museums visited previous year	3.31	4.75	3.54	3.86
GPA (0-10) ^a	6.73	6.82	6.98	6.84
Interest in sciences (1-3) ^b	2.12	1.74	2.23	2.03
Parental education (1-5) ^c	3.26	3.54	3.56	3.45
<i>FRIENDS' BACKGROUND CHARACTERISTICS</i>				
Male (1/0)	0.18	0.28	0.49	0.32
Already visited Palazzo Vecchio (1/0)	0.62	0.74	0.69	0.68
No. of museums visited previous year	3.07	4.59	3.20	3.61
GPA (0-10) ^a	6.66	6.78	6.61	6.68
Interest in sciences (1-3) ^b	2.09	1.68	2.06	1.95
Parental education (1-5) ^c	3.23	3.59	3.44	3.42
<i>STRUCTURE OF FRIENDSHIP NETWORKS</i>				
No. of friends (individual degree)	4.94	4.00	3.07	4.00
No. of friends of friends (friends' degree)	5.62	4.81	3.43	4.61
No. of classes	5	5	5	15
No. of students per class	17.8	17.4	18	17.7
No. of students	89	87	90	266

Note. Estimates are based on sampling-theory estimation of a ratio for cluster sampling (Cochran, 1963, pag. 30).

^a Grade Point Average (GPA) is continuous on a 1-10 point scale.

^b Levels of interest in sciences: 1 if the interest is only in humanities, 2 if it is in both humanities and sciences, 3 if it is only in sciences.

^c Levels of parental education: 1 if both parents completed only compulsory education, 2 if only one of them completed high school, 3 if both completed high school, 4 if only one of them completed university, 5 if both parents completed university.

always well balanced across encouragement groups. Imbalances between groups of students in classes assigned to the alternative encouragement levels arise in almost all the individual-level and network(friends)-level background characteristics. In particular, number of museums visited the previous year and proportion of males differ remarkably across students in classes exposed to alternative encouragement levels. Students assigned to presentation visited more museums, on average, than those assigned to the other encouragements; not surprisingly, they are also more interested in the humanities. Furthermore, the proportion of males is much higher among students assigned to the reward. Imbalances also arise in the structure of friendship networks. Specifically, the average number of ties (i.e., friends) is about 5, 4 and 3 in the flyer, presentation and reward group, respectively. These differences result in a different number of friends of friends across encouragement groups, too: students assigned to the reward tend to have less friends – and less friends of friends – than those assigned to the presentation and, to a greater extent, than those assigned to the flyer.

Table 2 reports descriptive statistics regarding treatment uptake (Palazzo Vecchio visit) and final outcome (number of museum visits during the follow-up), again at both individual/student and network/student’s friends level. In addition, Figure 1 shows the bar charts of the final outcome by encouragement status. The percentage of students who visit Palazzo

Table 2


The museum study: descriptive statistics of the treatment and the outcome variables

	<i>Flyer</i>	<i>Presentation</i>	<i>Reward</i>	<i>Overall</i>
<i>INDIVIDUAL TREATMENT AND OUTCOME</i>				
No. of students performing the proposed visit	3	10	40	53
Share of students performing the proposed visit	0.03	0.11	0.44	0.20
Count of museum visits during the follow-up	1.49	4.39	3.00	2.95
Students visiting at least one museum during the follow-up	0.36	0.86	0.91	0.71
<i>FRIENDS’ TREATMENT AND OUTCOME</i>				
Share of friends performing the proposed visit	0.03	0.10	0.41	0.18
Count of museum visits during the follow-up	1.57	4.02	2.86	2.80
Friends visiting at least one museum during the follow-up	0.56	0.98	0.92	0.82

Vecchio after receiving the flyer is only 3%. Among students who also attend the presenta-

Figure 1

Outcome bar charts



tion made by the art expert, this percentage increases to 11%. A sharper increase is observed when the reward is also promised in the event of a visit to Palazzo Vecchio: in this case, in fact, 44% of students performed the proposed visit. However, it is worth noting that the interest of the field experiment is not on investigating which incentive could better motivate students to perform a particular museum visit once, but on assessing which incentive is more effective to prompt teens' museums attendance in general. Simple comparisons of the outcomes by encouragement status suggest that the flyer is insufficient to motivate students to visit art museums in the follow-up period, just as it is insufficient to motivate students to visit Palazzo Vecchio in the first place. As we can see in Figure 1, the empirical distributions of the number of museum visits during the follow up by encouragement are left-skewed and have a mass at zero. The proportion of students who do not visit any museum is especially high under the flyer encouragement. The number of museum visits has a similar distribution under the presentation and reward encouragements, which seem to be more effective than the flyer encouragement. The average number of total museum visits is relatively high under both the presentation and the reward (4.39 and 3.00, respectively). Therefore, there is some evidence that these two types of encouragement may have at least partially achieved the purpose of motivating students to visit art museums, regardless of whether the triggered mechanism is ascribable to the encouragement itself, to the treatment or to spillovers from friends.

3 Methodology

3.1 Potential outcomes

Let $\mathbf{z}$ denote a J -dimensional vector of the encouragement assignments for the whole sample of classes with j th element equal to z_j : $\mathbf{z} \in \{1, 2, 3\}^J$. For each student i in class j , $i = 1 \dots, N_j$, $j = 1, \dots, J$, let $M_{ij}(\mathbf{z})$ denote the potential indicator for visiting Palazzo Vecchio given the encouragement vector $\mathbf{z}$; also let $Y_{ij}(\mathbf{z})$ denote the potential outcome of main interest, that is, the number of museums visited in the follow-up period given the encouragement vector $\mathbf{z}$. Because $S_{\mathcal{N}_{ij}}$ is function of $M_{i'j}$, $i' \neq i \in \mathcal{N}_{ij}$, it is a post-encouragement variable, and thus, we introduce potential outcomes for it, too:

$$S_{\mathcal{N}_{ij}}(\mathbf{z}) = \frac{1}{|\mathcal{N}_{ij}|} \sum_{i' \neq i \in \mathcal{N}_{ij}} M_{i'j}(\mathbf{z}).$$

We assume that a cluster-level Stable Unit Value Assumption (SUTVA, Rubin, 1980) holds for the potential outcomes $M_{ij}(\mathbf{z})$, $Y_{ij}(\mathbf{z})$, and $S_{\mathcal{N}_{ij}}(\mathbf{z})$:

Assumption 1. (SUTVA for $M_{ij}(\mathbf{z})$, $Y_{ij}(\mathbf{z})$, and $S_{\mathcal{N}_{ij}}(\mathbf{z})$)

No hidden variation of encouragement: For all $\mathbf{z}, \mathbf{z}'$: $\mathbf{z} = \mathbf{z}'$, $M_{ij}(\mathbf{z}) = M_{ij}(\mathbf{z}')$, $Y_{ij}(\mathbf{z}) = Y_{ij}(\mathbf{z}')$, $S_{\mathcal{N}_{ij}}(\mathbf{z}) = S_{\mathcal{N}_{ij}}(\mathbf{z}')$.

Cluster-level no-interference: For all $\mathbf{z}, \mathbf{z}'$: $z_j = z'_j$, $M_{ij}(\mathbf{z}) = M_{ij}(\mathbf{z}')$, $Y_{ij}(\mathbf{z}) = Y_{ij}(\mathbf{z}')$, $S_{\mathcal{N}_{ij}}(\mathbf{z}) = S_{\mathcal{N}_{ij}}(\mathbf{z}')$.

Assumption 1 allows to write $M_{ij}(\mathbf{z})$, $Y_{ij}(\mathbf{z})$ and $S_{\mathcal{N}_{ij}}(\mathbf{z})$ as $M_{ij}(z_j)$, $Y_{ij}(z_j)$ and $S_{\mathcal{N}_{ij}}(z_j) = |\mathcal{N}_{ij}|^{-1} \sum_{i' \neq i \in \mathcal{N}_{ij}} M_{i'j}(z_j)$. Assumption 1 implies that the outcomes of a given student does not vary with the encouragement assigned to students of other classes. Under SUTVA, $M_{ij} = M_{ij}(Z_j)$, $S_{ij} = S_{ij}(Z_j)$, and $Y_{ij} = Y_{ij}(Z_j)$.

Our key insight is that spillover effects may pass through the variable $S_{\mathcal{N}_{ij}}$. We formalize this intuition by viewing it as a mediator. For $j = 1, \dots, J$, let $\mathbf{S}_j$ and $\mathbf{S}_j(z_j)$ denote N_j -dimensional vectors with i th element equal to $S_{\mathcal{N}_{ij}}$ and $S_{\mathcal{N}_{ij}}(z_j)$, respectively, $i = 1 \dots, N_j$. For each student i , we assume the existence of potential outcomes of the form $Y_{ij}(z_j, \mathbf{s}_j)$ and $Y_{ij}(z_j, \mathbf{S}_j(z'_j))$: $Y_{ij}(z_j, \mathbf{s}_j)$ would be the value of the outcome Y if, possibly

contrary to fact, the encouragement were set to the level z_j and the mediator vector, $\mathbf{S}_j$, were set to a specific prefixed value, $\mathbf{s}_j$; and $Y_{ij}(z_j, \mathbf{S}_j(z'_j))$ would be the value of the outcome Y if, possibly contrary to fact, the encouragement were set to the level z_j and the mediator vector, $\mathbf{S}_j$, were set to the value it would have taken if the encouragement had been set to an alternative level, z'_j .

We believe to be plausible that a cluster version of the *Stable Unit Treatment on Neighborhood Value Assumption* (SUTNVA) introduced by Forastiere, Airoidi, & Mealli (2021) holds for potential outcomes of the form $Y_{ij}(z_j, \mathbf{s}_j)$ and $Y_{ij}(z_j, \mathbf{S}_j(z'_j))$. Formally,

Assumption 2. (Cluster SUTNVA)

No hidden variation of mediators:

For all $\mathbf{s}_j, \mathbf{s}'_j$: $\mathbf{s}_j = \mathbf{s}'_j$, $Y_{ij}(z_j, \mathbf{s}_j) = Y_{ij}(z_j, \mathbf{s}'_j)$.

For all $\mathbf{S}_j(z'_j), \mathbf{S}_j(z''_j)$: $\mathbf{S}_j(z'_j) = \mathbf{S}_j(z''_j)$, $Y_{ij}(z_j, \mathbf{S}_j(z'_j)) = Y_{ij}(z_j, \mathbf{S}_j(z''_j))$.

Cluster Neighborhood Interference:

For all $\mathbf{s}_j, \mathbf{s}'_j$: $\mathbf{s}_{\mathcal{N}_{ij}} = \mathbf{s}'_{\mathcal{N}_{ij}}$, $Y_{ij}(z_j, \mathbf{s}_j) = Y_{ij}(z_j, \mathbf{s}'_j)$.

For all $\mathbf{S}_j(z'_j), \mathbf{S}_j(z''_j)$: $\mathbf{S}_{\mathcal{N}_{ij}}(z'_j) = \mathbf{S}_{\mathcal{N}_{ij}}(z''_j)$, $Y_{ij}(z_j, \mathbf{S}_j(z'_j)) = Y_{ij}(z_j, \mathbf{S}_j(z''_j))$.

Cluster SUTNVA implies that potential outcomes of the form $Y_{ij}(z_j, \mathbf{s}_j)$ and $Y_{ij}(z_j, \mathbf{S}_j(z'_j))$ for student i in class j depend on proportion of student i 's friends who visit Palazzo Vecchio, but do not vary with the proportion of friends of other classmate students who visit Palazzo Vecchio. Under cluster SUTNVA (Assumption 2), we can write $Y_{ij}(z_j, \mathbf{s}_j)$ and $Y_{ij}(z_j, \mathbf{S}_j(z'_j))$ as $Y_{ij}(z_j, \mathbf{s}_{\mathcal{N}_{ij}})$ and $Y_{ij}(z_j, \mathbf{S}_{\mathcal{N}_{ij}}(z'_j))$, respectively. Note that, under Assumptions 1 and 2, $Y_{ij} = Y_{ij}(Z_j) = Y_{ij}(Z_j, \mathbf{S}_{\mathcal{N}_{ij}}(Z_j))$.

It is worth noting that we have to know the network that links the units to exploit Assumptions 1 and 2, which are based on the idea that interference can take place between immediate neighbors. Moreover we need to assume that this network is fixed and that the links between nodes are correctly measured. In our setting this assumption means that any change in the vector of assignment $\mathbf{z}$ and in $\{\mathbf{S}_j(z_j)\}_{j=1}^J$ does not modify the network that links the subjects.

In the causal inference literature, potential outcomes of the form $Y_{ij}(z_j, s_{\mathcal{N}_{ij}})$ and $Y_{ij}(z_j, S_{\mathcal{N}_{ij}}(z'_j))$ are also called *a priori* counterfactual potential outcomes (e.g., Frangakis & Rubin, 2002), because some of them can never be observed for some units. For instance, $Y_{ij}(z_j, S_{\mathcal{N}_{ij}}(z'_j))$ cannot be observed for any units with $S_{\mathcal{N}_{ij}}(z'_j) \neq S_{\mathcal{N}_{ij}}(z_j)$. In the sequel, we denote by $M_{ij}(z)$, $S_{\mathcal{N}_{ij}}(z)$ and $Y_{ij}(z)$ potential outcomes for M_{ij} , $S_{\mathcal{N}_{ij}}$ and Y_{ij} for a student i in class j under class encouragement z . Similarly, we use the notation $Y_{ij}(z, s_{\mathcal{N}_{ij}})$ and $Y_{ij}(z, S_{\mathcal{N}_{ij}}(z'))$ for class encouragements z and z' .

3.2 The Principal Stratification Approach

In clustered encouragement designs with individual noncompliance, we are often interested in causal effects defined accounting for the compliance behavior. We deal with this issue using the principal stratification framework.

Principal stratification was proposed by Frangakis & Rubin (2002) as a general framework for the evaluation of treatment effects while addressing complications arising after the assignment to an encouragement or a treatment, like non-compliance (e.g., Hirano et al., 2000; Frangakis et al., 2002; Forastiere et al., 2016) and truncation by death (e.g., Mattei & Mealli, 2007). A principal stratification with respect to a post-assignment variable is a classification of units into latent groups, called “principal strata”, defined by the joint potential values of that post-assignment variable under alternative levels of the assignment. Principal strata are latent groups because it is generally not possible to observe all the potential values of the post-assignment variable for any unit: it is only possible to observe the potential outcomes associated with the actual assignment, while the potential outcome under any alternative assignment is missing. Therefore, the specific stratum a unit belongs to is generally unknown.

In our study, we classify students into principal strata defined by their compliance behavior. Because the encouragement is a three-valued variable, $Z_j \in \{1, 2, 3\}$, and the intermediate variable is a binary indicator for the visit Palazzo Vecchio by two months after

randomization of the encouragement, $M_{ij} \in \{0, 1\}$, the principal stratification with respect to M_{ij} classifies units into eight latent subgroups according to the joint potential values of M_{ij} under each encouragement level: $M_{ij}(1)$, $M_{ij}(2)$ and $M_{ij}(3)$. More formally, let $G_{ij} = (M_{ij}(1), M_{ij}(2), M_{ij}(3))$, then the principal strata are given by the possible values of G_{ij} : $(m_1, m_2, m_3) := \{ij : M_{ij}(1) = m_1, M_{ij}(2) = m_2, M_{ij}(3) = m_3\}$ with $m_1, m_2, m_3 \in \{0, 1\}$. We relabel the possible values of G_{ij} as $m_1m_2m_3$, so $G_{ij} \in \{111, 011, 101, 110, 001, 010, 100, 000\}$. We name the eight principal strata as shown in Table 3.

Table 3

Principal strata and their labels

$M_{ij}(1)$	$M_{ij}(2)$	$M_{ij}(3)$	<i>Principal Stratum</i>	<i>Stratum Label</i>
1	1	1	$G = 111$	<i>Always Takers</i>
0	1	1	$G = 011$	<i>Presentation Compliers</i>
0	0	1	$G = 001$	<i>Reward Compliers</i>
0	0	0	$G = 000$	<i>Never Takers</i>
1	0	0	$G = 100$	<i>Defiers</i>
0	1	0	$G = 010$	<i>Defiers</i>
1	1	0	$G = 110$	<i>Defiers</i>
1	0	1	$G = 101$	<i>Defiers</i>

As we can see in Table 3, we refer to the union of principal strata 100, 010, 110, 101 as *Defiers*. They are students who would visit Palazzo Vecchio under the flyer and/or presentation condition(s), but would not visit the museum under a stronger encouragement (presentation and/or reward). In our application study, similar behaviors are implausible: if the flyer and/or the presentation is/are enough to motivate a student to visit Palazzo Vecchio, we see no reasons why receiving additional motivators should discourage the same visit. Therefore, we rule out the presence of Defiers making the following assumption:

Assumption 3. (Monotonicity of Compliance). For all students i in each class j , $M_{ij}(1) \leq M_{ij}(2) \leq M_{ij}(3)$.

Under monotonicity, $G_{ij} \in \{111, 011, 001, 000\}$ for all ij , that is, each student ij can belong to one of the following four groups: *Always Takers* (AT), *Presentation Compliers* (PC),

Reward Compliers (RC) and *Never Takers* (NT). *Always Takers* and *Never Takers* are students who, respectively, would perform, and would not perform the proposed visit to Palazzo Vecchio regardless of the encouragement. *Presentation Compliers* are students who would visit Palazzo Vecchio if they received the reward or at least attended the presentation. *Reward Compliers* are students who would visit Palazzo Vecchio only upon receiving the reward.

3.3 Principal Causal Effects

The main advantage of the principal stratification approach is that it allows to define stratum-specific causal effects, named *Principal Causal Effects* (PCEs), solving the complication of adjusting causal effects for the post-assignment variable, M_{ij} , the actual visit to Palazzo Vecchio. Because principal strata are not affected by encouragement assignment, PCEs are properly defined causal effects.

We focus on finite population average principal causal effects. Overall average principal causal effects of an encouragement condition, say z , versus another, say z' , are defined as follows:

$$PCE_g(z \text{ vs } z') = \frac{\sum_{ij:G_{ij}=g} Y_{ij}(z) - Y_{ij}(z')}{N_g} = \frac{\sum_{ij:G_{ij}=g} Y_{ij}(z, S_{\mathcal{N}_{ij}}(z)) - Y_{ij}(z', S_{\mathcal{N}_{ij}}(z'))}{N_g} \quad (1)$$

where N_g is the total number of subjects in the sample belonging to the principal stratum $g \in \{000, 001, 011, 111\}$.

For Presentation and Reward Compliers, overall principal causal effects are *associative*, in the sense that they measure effects on the outcome that are associative with effects on the intermediate variable, M_{ij} , combining three effects: the *pure encouragement effect*, the *experience effect*, and the *spillover effect* (Frangakis & Rubin, 2002; Forastiere, Lattarulo, et al., 2021). Pure encouragement effects measure changes in the individual primary outcome, i.e., number of museum visits in the follow-up period, directly induced by the encouragement.

Experience effects measure changes in the individual outcome that originate from individual treatment uptake, i.e., from visiting Palazzo Vecchio first-hand. Finally, spillover effects measure changes in the individual outcome ascribable to mechanisms of interference sparked by friends performing the encouraged visit to Palazzo Vecchio in the first place. Spillover effects can be interpreted as channeled effects: they measure the part of the encouragement effects mediated through the proportion of friends visiting Palazzo Vecchio.

For Never Takers and Always Takers, overall principal causal effects are *dissociative* (Frangakis & Rubin, 2002): because the visit to Palazzo Vecchio is – by definition – not affected by the encouragement for Never Takers and Always Takers, it follows that any change in the individual outcome may be “only” ascribed to pure encouragement and spillover effects (Forastiere, Lattarulo, et al., 2021).

In Table 4, for each principal stratum, we resume the different causal pathways involved in the *PCEs*. As we can see in Table 4, we interpret the overall principal causal effects of

Table 4

Causal mechanisms underlying Principal Causal Effects

<i>Principal Stratum</i>	<i>PCE_g(3 vs 2)</i>	<i>PCE_g(2 vs 1)</i>	<i>PCE_g(3 vs 1)</i>
Never Takers	Spill	Enc, Spill	Enc, Spill
Always Takers	Spill	Enc, Spill	Enc, Spill
Presentation Compliers	Spill	Enc, Exp, Spill	Enc, Exp, Spill
Reward Compliers	Exp, Spill	Enc, Spill	Enc, Exp, Spill

Note. “Enc”, “Exp” and “Spill” stand for, respectively, pure encouragement effect, experience effect and spillover effect.

reward ($z = 3$) versus presentation ($z = 2$) for Never Takers, Always Takers and Presentation Compliers as spillover effects, implicitly assuming that there do not exist pure encouragement and experience effects for these types of subjects. Indeed, for Never Takers ($G_{ij} = 000$), Always Takers ($G_{ij} = 111$), and Presentation Compliers ($G_{ij} = 011$), $M_{ij}(2) = M_{ij}(3)$, and thus, experience effects do not exist by definition. Moreover, for these types of students, we can reasonably assume that there exist no pure encouragement effect of reward ($z = 3$)

versus presentation ($z = 2$). Under assignment to the reward, students get the reward only if they visit Palazzo Vecchio within two months after assignment; they do not get any reward if they visit art museums in the follow-up period but do not visit Palazzo Vecchio in the first place. Therefore, in terms of pure encouragement effects we can argue that the reward adds nothing with respect to presentation.

3.4 Disentangling Spillovers from PCEs Through Mediation

This section describes the core of our novel methodological proposal: using principal stratum mediation analysis to disentangle spillovers from the previous PCEs.

In principal stratum mediation analysis, focus is on disentangling the part of each principal causal effect that works through another variable, a mediator – *indirect principal causal effect*, from the causal effects of the assignment/treatment on the response net of the effect of the mediator – *direct principal causal effect*. In our application study, we conduct principal stratum mediation analysis using the proportion of friends who visited Palazzo Vecchio as a mediator. Therefore, indirect principal causal effects measure changes in the response that can be ascribed to a variation in the summary of friends' compliance behavior by principal strata, and thus, they are interpretable as spillover effects. Formally, we introduce the concepts of principal natural direct and indirect effects and principal control direct effects:

$$NDE_g(z \text{ vs } z', S_{\mathcal{N}_{ij}}(z'')) = \frac{\sum_{ij:G_{ij}=g} [Y_{ij}(z, S_{\mathcal{N}_{ij}}(z'')) - Y_{ij}(z', S_{\mathcal{N}_{ij}}(z''))]}{N_g} \quad (2)$$

$$NIE_g(z'', S_{\mathcal{N}_{ij}}(z) \text{ vs } S_{\mathcal{N}_{ij}}(z')) = \frac{\sum_{ij:G_{ij}=g} [Y_{ij}(z'', S_{\mathcal{N}_{ij}}(z)) - Y_{ij}(z'', S_{\mathcal{N}_{ij}}(z'))]}{N_g} \quad (3)$$

$$CDE_g(z \text{ vs } z', s_{\mathcal{N}_{ij}}^*) = \frac{\sum_{ij:G_{ij}=g} [Y_{ij}(z, s_{\mathcal{N}_{ij}}^*) - Y_{ij}(z', s_{\mathcal{N}_{ij}}^*)]}{N_g} \quad (4)$$

for $z, z', z'' \in \{1, 2, 3\}$ with $z \neq z'$ and z'' set to z or to z' , and $s_{\mathcal{N}_{ij}}^* \in [0, 1]$ is a value of the mediator reasonably chosen by the researcher and not necessarily observed in the sample.

In the application, we calculate these estimands for $z > z'$.

Principal natural direct effects measure average local effects (that is, for students of type g) of z versus z' on the response, the number of visits to art museums, not mediated through $S_{\mathcal{N}_{ij}}$, the proportion of friends visiting Palazzo Vecchio; in contrast, they consider the effect of z versus z' on the response intervening by fixing the proportion of friends visiting Palazzo Vecchio to the value it would have taken if the encouragement had been set to z'' . Principal natural direct effects embrace all causal channels other than spillovers, i.e., the pure encouragement effect and, where appropriate, also experience effects.

Principal natural indirect effects fix the encouragement to the value z'' and then measure average local effects (that is, for students of type g) on the response, the number of visits to art museums, mediated through the proportion of friends visiting Palazzo Vecchio, $S_{\mathcal{N}_{ij}}$; that is, intervening by setting $S_{\mathcal{N}_{ij}}$ to what it would have been if the encouragement were z , $S_{\mathcal{N}_{ij}}(z)$, in contrast to what it would have been if the encouragement were z' , $S_{\mathcal{N}_{ij}}(z')$. Principal natural indirect effects provide information on spillover effects.

Principal controlled direct effects measure local effects on the response not mediated through the proportion of friends visiting Palazzo Vecchio, as principal natural direct effects, but they compare potential outcomes for students belonging to the same principal stratum under two encouragements, z and z' , setting the mediator, $S_{\mathcal{N}_{ij}}$, to some pre-fixed value $s_{\mathcal{N}_{ij}}^*$.

It is worth noting that “Natural Direct Effects” (NDE), “Natural Indirect Effects” (NIE), and “Controlled Direct Effects” (CDE), defined for the whole population, rather than for latent sub-populations such as the principal strata, are the causal estimands usually considered in mediation analysis (e.g., Robins & Greenland, 1992; Pearl, 2001).

Note that, as in the classical mediation setting overall average causal effects can be decomposed into the sum of natural direct and indirect effects, each average overall principal causal effect can be decomposed into the sum of a principal natural direct effect and a principal natural indirect effect: $PCE_g(z \text{ vs } z') = NDE_g(z \text{ vs } z', S_{\mathcal{N}_{ij}}(z)) + NIE_g(z', S_{\mathcal{N}_{ij}}(z') \text{ vs } S_{\mathcal{N}_{ij}}(z)) = NDE_g(z' \text{ vs } z, S_{\mathcal{N}_{ij}}(z')) + NIE_g(z, S_{\mathcal{N}_{ij}}(z) \text{ vs } S_{\mathcal{N}_{ij}}(z'))$.

We also introduce new causal estimands, which we call principal controlled spillover effects, defined as follows:

$$CSE_g(z, s_{\mathcal{N}_{ij}}^{***} \text{ vs } s_{\mathcal{N}_{ij}}^{**}) = \frac{\sum_{ij:G_{ij}=g} [Y_{ij}(z, s_{\mathcal{N}_{ij}}^{***}) - Y_{ij}(z, s_{\mathcal{N}_{ij}}^{**})]}{N_g} \quad (5)$$

where $z \in \{1, 2, 3\}$ and $s_{\mathcal{N}_{ij}}^{**}$ and $s_{\mathcal{N}_{ij}}^{***} \in [0, 1]$ are values of the proportion of friends visiting Palazzo Vecchio. Principal controlled spillover effects measure local effects on the response for students belonging to the same principal stratum and exposed to the encouragement level z , due to a change in the proportion of friends visiting Palazzo Vecchio from a pre-fixed value $s_{\mathcal{N}_{ij}}^{***}$ to another pre-fixed value $s_{\mathcal{N}_{ij}}^{**}$. Therefore, principal controlled spillover effects provide information on spillover effects.

As discussed in Section 3.3, experience effects cannot exist for Always Takers and Never Takers. Therefore, for these two principal strata, principal natural and controlled direct effects provide information on pure encouragement effects, and principal natural indirect effects and principal controlled spillover effects provide information on spillover effects.

For Presentation Compliers and Reward Compliers, in general, there exist experience effects, and thus, overall principal causal effects for these principal strata combine encouragement, experience and spillover effects. Therefore, for Presentation Compliers and Reward Compliers, principal natural indirect effects and principal controlled spillover effects are interpretable as spillover effects, and principal natural and controlled direct effects entangle pure encouragement and experience effects. Table 5 summarizes the information on the causal pathways provided by principal natural and controlled direct effects. Note that there is no encouragement and experience effect for Never Takers, Always Takers and Presentation Compliers when encouragement levels 2 and 3 are compared, and thus, $NDE_g(3 \text{ vs } 2, S_{\mathcal{N}_{ij}}(z''))$ and $CDE_g(3 \text{ vs } 2, s_{\mathcal{N}_{ij}}^*)$ for $g \in \{000, 111, 011\}$ are zero.

Table 5*Principal natural direct effects and principal controlled direct effects: an interpretation*

<i>Principal Stratum</i> (g)	$NDE_g(3 \text{ vs } 2, S_{N_{ij}}(z''))$ $CDE_g(3 \text{ vs } 2, s_{N_{ij}}^*)$	$NDE_g(2 \text{ vs } 1, S_{N_{ij}}(z''))$ $CDE_g(2 \text{ vs } 1, s_{N_{ij}}^*)$	$NDE_g(3 \text{ vs } 1, S_{N_{ij}}(z''))$ $CDE_g(3 \text{ vs } 1, s_{N_{ij}}^*)$
Never Takers	0	Enc	Enc
Always Takers	0	Enc	Enc
Presentation Compliers	0	Enc, Exp	Enc, Exp
Reward Compliers	Exp	Enc	Enc, Exp

Note. “Enc” and “Exp” stand for pure encouragement effect experience effect, respectively.

3.5 Ignorability Assumptions

Our study is a CED, and thus, the encouragement is randomly assigned by design, implying that the encouragement is independent of all the potential outcomes and covariates. Formally,

Assumption 4. (Randomization of the encouragement)

$$Z_j \perp\!\!\!\perp \{\{M_{ij}(z)\}_z, \{S_{N_{ij}}(z)\}_z, \{Y_{ij}(z, s_{N_{ij}})\}_{z, s_{N_{ij}}}, \mathbf{X}_{ij}\} \text{ for } z \in \{1, 2, 3\} \text{ and } s_{N_{ij}} \in [0, 1]$$

Under monotonicity (Assumption 3), there exist four principal strata and most observed groups defined by the encouragement actually received and the observed value of M_{ij} host mixtures of principal strata as shown in Table 6.

Table 6

Observed groups of students and possible principal strata: NT=000, RC=001, PC=011, AT=111

Z_j	M_{ij}	No.	$\% Z_j = z$	$Mean(Y_{ij})$	$P(Y_{ij} > 0)$	<i>Possible principal Strata</i>
1	0	86	96.6%	1.35	0.34	PC, RC, NT
1	1	3	3.4%	5.67	1.00	AT
2	0	77	86.5%	4.58	0.87	RC, NT
2	1	10	11.2%	2.90	0.80	AT, PC
3	0	50	55.6%	3.18	0.88	NT
3	1	40	44.4%	2.78	0.95	AT, PC, RC

Randomization guarantees that principal stratum membership has the same distribution in all encouragement arms. Therefore, under randomization and monotonicity, we can

non-parametrically point identify principal stratum proportions. Let $\pi_{m_1 m_2 m_3}$ denote the proportion of units belonging to principal stratum $m_1 m_2 m_3 = \{ij : M_{ij}(1) = m_1, M_{ij}(2) = m_2, M_{ij}(3) = m_3\}$, that is, $\pi_{m_1 m_2 m_3} := P(G_{ij} = m_1 m_2 m_3)$. We have $\pi_{100} = \pi_{010} = \pi_{110} = \pi_{101} = 0$; $\pi_{111} = P(M_{ij} = 1 \mid Z_j = 1)$; $\pi_{000} = P(M_{ij} = 0 \mid Z_j = 3)$; $\pi_{001} = P(M_{ij} = 0 \mid Z_j = 2) - \pi_{000} = P(M_{ij} = 1 \mid Z_j = 3) - P(M_{ij} = 1 \mid Z_j = 2)$; and $\pi_{011} = P(M_{ij} = 1 \mid Z_j = 2) - \pi_{111} = P(M_{ij} = 1 \mid Z_j = 3) - \pi_{111} - \pi_{001}$. The method-of-moments estimates of the principal stratum proportions and also Table 6) are as follows (standard errors in parenthesis): $\hat{\pi}_{111} = 0.034$ (0.019); $\hat{\pi}_{000} = 0.556$ (0.052); $\hat{\pi}_{001} = 0.329$ (0.062); $\hat{\pi}_{011} = 0.081$ (0.039); suggesting that there are more than 50% Never Takers, and less than 3.5% Always Takers. About 1/3 of students are expected to be Reward Compliers and 8.1% of students are expected to be Presentation Compliers.

Although we can point identify principal stratum proportions, principal stratum membership is only partially observed, and thus, overall principal causal effects, $PCE_g(z \text{ vs } z')$, are not non-parametrically point identified without additional assumptions. Inference on principal natural direct and indirect effects (Equations 2 and 3) and principal controlled direct and spillover effects (Equations 4 and 5) is even more challenging, because these estimands involve potential outcomes of the form $Y_{ij}(z, s_{\mathcal{N}_{ij}})$ and $Y_{ij}(z, S_{\mathcal{N}_{ij}}(z'))$, which are never observed in this specific experiment for some students. Some additional structural assumptions that allow us to extrapolate information on this type of potential outcomes from at least potentially observable data is required. Here we invoke “latent” ignorability assumptions of the mediator, that is, ignorability assumptions of the mediator conditional on principal stratum membership.

Assumption 5. (Latent ignorability of the mediator (I)). For each unit i in cluster j , for all $z \in \{1, 2, 3\}$ and $s_{\mathcal{N}_{ij}} \in [0, 1]$, $Y_{ij}(z, s_{\mathcal{N}_{ij}}) \perp\!\!\!\perp S_{\mathcal{N}_{ij}} \mid Z_j, G_{ij}, \mathbf{X}_{ij}$;

Assumption 6. (Latent ignorability of the mediator (II)). For each unit i in cluster j , for all $z, z' \in \{1, 2, 3\}$ and $s_{\mathcal{N}_{ij}} \in [0, 1]$, $Y_{ij}(z, s_{\mathcal{N}_{ij}}) \perp\!\!\!\perp S_{\mathcal{N}_{ij}}(z') \mid G_{ij}, \mathbf{X}_{ij}$.

In our application study, Assumption 5 means that potential outcomes of the form

$Y_{ij}(z, s_{\mathcal{N}_{ij}})$ for future visits to other museums of a student i in class j are independent of the decision to visit Palazzo Vecchio of subject i 's friends in class j conditional on student i 's encouragement status, covariates and compliance status. Assumption 6 implies that $Y_{ij}(z, s_{\mathcal{N}_{ij}})$ and the decision to visit Palazzo Vecchio of subject i 's friends in class j under encouragement z' , $S_{\mathcal{N}_{ij}}(z')$, are independent conditional on student i 's covariates and compliance status.

If the compliance status were fully observed for each unit, Assumption 5 and Assumption 6 would be the standard ignorability assumptions on the mediator usually invoked in the causal mediation literature (see, e.g., VanderWeele & Vansteelandt, 2009; Ten Have & Joffe, 2012, for a review). Therefore, we could identify principal controlled direct and spillover effects under Assumptions 4 and 5 and principal natural direct and indirect effects under Assumptions 4, 5 and 6 (see the online Appendix for details). Because the compliance status is only partially observed, Assumptions 4, 5 and 6 do not solve identification issues. We deal with inferential issues arising from the latent nature of principal stratum membership using a model-based Bayesian approach, which does not need full identification (Imbens & Rubin, 1997). A Bayesian model-based principal stratification analysis requires the specification of two sets of models: one for the principal strata membership given the covariates, and one for the outcomes conditional on the covariates and the principal strata. In addition, a prior distribution for the model parameters must be specified.

4 Analyses

4.1 Parametric Models

In our setting, randomization is at cluster-level, and thus, a within-class correlation among individuals may arise from reciprocal influence or from other unobserved common factors. In particular, students belonging to the same class are likely to show resemblance in terms of both compliance behavior and final outcomes. We account for this systematic unexplained

variation within classes introducing a hierarchical structure, by specifying varying-intercept models for the compliance status and for the outcome.

Principal Strata Model

We specify a multinomial logit random effect model for principal stratum membership conditional on covariates:

$$\log \left(\frac{P(G_{ij} = g | \mathbf{X}_{ij})}{P(G_{ij} = 111 | \mathbf{X}_{ij})} \right) = \gamma_g + \boldsymbol{\delta}_g^T \mathbf{X}_{ij} + a_j \quad \text{with } a_j \sim N(0, \sigma_a), \quad (6)$$

where $g \in \{000, 001, 011\}$ (the stratum 111 is chosen as the reference category) and a_j are the random intercepts.

The vector $\mathbf{X}_{ij}$ includes the individual background characteristics introduced in Table 1: *Male* (1/0), *Already visited Palazzo Vecchio* (1/0), *No. of museums visited previous year*, *GPA* (continuous on a 1-10 points scale), *Interest in human sciences* (1/0) and *Parental education* (1 if at least one of the parents completed university, 0 otherwise). In addition, it includes a network variable indicating whether friends already visited Palazzo Vecchio prior to the experiment or not.

Potential Outcome Models

We specify Zero-Inflated Poisson (ZIP) models for the potential outcomes conditional on principal stratum membership and covariates. A ZIP model assumes that subjects are a mixture of two types. For subjects of first type the outcome is always zero. For subjects of the second type, the outcome follows an usual Poisson distribution, which can produce the zero outcome or some other. Let $\phi_{ij,g,z}$ denote the conditional probability given the covariates that an individual with compliance behaviour g under assignment z is of the first type, and let $\mu_{ij,g,z}$ denote the Poisson mean, which we model using a hierarchical generalized

linear model with a log link:

$$Pr(Y_{ij}(z, S_{\mathcal{N}_{ij}}(z)) = y \mid G_{ij} = g, \mathbf{X}_{ij}) = \begin{cases} \phi_{ij,g,z} + (1 - \phi_{ij,g,z}) \cdot Pois(0; \mu_{ij,g,z}) & \text{if } y = 0 \\ (1 - \phi_{ij,g,z}) \cdot Pois(y; \mu_{ij,g,z}) & \text{if } y > 0 \end{cases} \quad (7)$$

where

$$Pois(y; \mu_{ij,g,z}) = \frac{e^{-\mu_{ij,g,z}} \mu_{ij,g,z}^y}{y!} \quad y = 0, 1, 2, \dots$$

In the ZIP models for the potential outcomes, we condition on the same individual and friends' background characteristics we include in the principal strata model and on an additional network variable indicating the number of friends' past museums visits. We include the latter because we can reasonably believe that having museum goers friends may influence the individual museums attendance.

As we discussed in Section 3.3, we can reasonably assume that there is no encouragement effects and experience effects for Never Takers, Always Takers and Presentation Compliers in the comparison of potential outcomes under reward ($z = 3$) versus presentation ($z = 2$) (see Table 5). This assumption implies that for $g \in \{000, 111, 011\}$, $\phi_{ij,g,2} = \phi_{ij,g,3}$ and $\mu_{ij,g,2} = \mu_{ij,g,3}$, which, in turn, implies that $\alpha_{g,2} = \alpha_{g,3}$, $\beta_{g,2}^{(S)} = \beta_{g,3}^{(S)}$ and $\beta_{g,2}^{(X)} = \beta_{g,3}^{(X)}$.

For the sake of parsimony, we assume that the probabilities of extra zeros with respect to the Poisson distribution, $\phi_{ij,g,z}$, depend on the encouragement and the principal stratum membership, but not on covariates: $\phi_{ij,g,z} = \phi_{g,z}$. Moreover, we impose the following prior equalities of the slope coefficients and the standard deviations of the random intercepts in the Poisson regressions: $\beta_{g,z}^{(S)} \equiv \beta_z^{(S)}$, for $g \in \{000, 001, 011, 111\}$ and $z \in \{1, 2\}$, $\beta_{g,z}^{(X)} \equiv \beta^{(X)}$ for all $z \in \{1, 2, 3\}$ and $g \in \{000, 001, 011, 111\}$, and $\sigma_{b,z} \equiv \sigma_b$, for all $z \in \{1, 2, 3\}$.

4.2 Bayesian Inference

We use a Bayesian approach for inference, under the assumption that parameters are a priori independent. We specify relative weakly informative prior distributions, using normal priors

for the coefficients of the principal strata model and of the potential outcome models and half-normal prior distributions for the standard deviations σ_a and σ_b of the random intercepts a_j and b_j entering the principal strata model and the potential outcome models (Gelman, 2006). The posterior distribution is approximated via Stan, a Bayesian programming language that implements a variant of Hamiltonian Monte Carlo (Stan Development Team, 2014), and by using a data augmentation approach. The latter involves, at each iteration, the imputation of the missing potential outcomes for Y_{ij} and M_{ij} for each unit. Then, the effects are estimated as function of the observed and missing information. See the online appendix for details on prior specifications, likelihood function and computational details.

5 Results

5.1 Principal Strata

Table 7 shows some summary statistics for the posterior distribution of the principal strata membership. The estimated proportions of the principal strata suggest there are more than 58% Never Takers, more than 30% Reward Compliers, less than 10% Always Takers and only 2% Presentation Compliers. Therefore, the majority of students do not visit Palazzo Vecchio irrespective of the encouragement, and a large minority of students is induced to visit Palazzo Vecchio if a reward is promised. Next to these two large strata, there are other more numerically negligible minorities for whom any encouragement works or for whom the presentation constitutes a sufficient motivator.

Table 7

Summary statistics of the posterior probabilities of principal stratum membership

	<i>Mean</i>	<i>SD</i>	<i>2.5%</i>	<i>5%</i>	<i>95%</i>	<i>97.5%</i>
Always Takers	0.095	0.025	0.053	0.056	0.135	0.143
Presentation Compliers	0.020	0.022	0.000	0.000	0.068	0.079
Reward Compliers	0.301	0.049	0.203	0.218	0.376	0.395
Never Takers	0.585	0.037	0.515	0.526	0.647	0.658


It is worth noting that the posterior means for the principal stratum proportions are slightly different from the method-of-moment estimates, especially for the proportion of Always Takers and Presentation Compliers. These differences are mainly attributable to the fact that the method-of-moment estimates do not account for individual and network level student's pre-experimental characteristics, which may affect the estimates given the observed imbalance across encouragement groups. Nevertheless, results are consistent, in the sense that both the unadjusted method-of-moment estimates and the adjusted model-based Bayesian estimates show evidence that the vast majority of students are either Never Takers or Reward Compliers and a small proportion of students are either Always Takers or Presentation Compliers.

An appealing advantage of the model-based Bayesian approach we use is that it easily allows us to investigate the distribution of baseline characteristics within each principal stratum, which may provide precious information on principal stratum members and on the propensity of members' friends to belong to the same latent principal stratum.

Figure 2 shows the boxplots of the posterior distributions of individual and friends' covariates in each principal stratum. All strata comprise students with similar GPAs, but they are somewhat different in terms of the other characteristics.

Figure 2

Distribution of covariates by principal strata



The small clutch of Always Takers is mostly composed of male students from fairly educated families, whose scholastic interests are slightly more in the sciences than in the humanities, and that visited some museums during their spare time in the pre-experimental period. They are very likely to have already visited Palazzo Vecchio. Also, their friends are quite likely to have visited Palazzo Vecchio previously.

The stratum of Never Takers is populated by a large majority of females, mostly from educated families, equally interested in the sciences and humanities. They have previously visited some museums in their free time. Both Never Takers and their friends are likely to have already visited Palazzo Vecchio.

The tiny stratum of Presentation Compliers mostly consists of males, from fairly educated families, with a marked preference for the sciences over the humanities. They are slightly less accustomed than others to spending their free time in museums. They are also less likely than the members of other strata, and also than their own friends, to have already visited Palazzo Vecchio.


Finally, the stratum of Reward Compliers mostly hosts female students from not very educated families, with a marked preference for the humanities over the sciences. They have

already dedicated their free time to visiting some museums, but they are slightly less likely than Never Takers and Always Takers to have already visited Palazzo Vecchio. Their friends are also slightly less likely to have visited Palazzo Vecchio previously than Never Takers' and Always Takers' friends.

Figure 3 shows the boxplots of the posterior distributions of the proportions of friends in each stratum for a student belonging to a given stratum, providing insights on the extent to which befriended students belong to the same principal stratum. In a sense, this kind of analysis is informative about the existence of homophily in terms of the latent characteristics of the befriended students rather than, as it is usually done in the literature, in terms of their observable characteristics. As shown in Figure 3, Always Takers are inclined to be befriended with other Always Takers and Never Takers, not so much with Reward Compliers and Presentation Compliers. Presentation Compliers are equally likely to be befriended with the members of all the other strata, but it is more difficult for them to be friends with their stratum peers. Reward Compliers tend to have relationships with Never Takers and, to a lesser extent, with their peers. Finally, Never Takers are most likely to have ties with their peers and, to a far lesser extent, with Reward Compliers.

Figure 3

Boxplots of the posterior distributions of proportion of friends in each stratum for a student belonging to a given stratum



5.2 PCEs and Principal Natural Direct and Indirect Effects

Of the four strata described above, we now focus on the two largest ones, i.e., Never Takers and Reward Compliers. For these two principal strata, Table 8 reports summary statistics of the posterior distributions of the finite-population overall principal causal effects and principal natural direct and indirect effects on count of museum visits during the follow-up period.

Table 8

Summary statistics of the posterior distributions of overall principal causal effects and principal natural direct and indirect effects for Never Takers and Reward Compliers

Causal Estimand	Never Takers ($G_{ij} = 000$)					Reward Compliers ($G_{ij} = 001$)				
	Mean	SD	5%	95%	Pr(>0)	Mean	SD	5%	95%	Pr(>0)
$PCE_g(2 \text{ vs } 1)$	1.65	1.04	0.07	3.35	0.96	4.73	1.42	2.66	7.14	1.00
$NDE_g(2 \text{ vs } 1, S_{N_{ij}}(2))$	1.66	1.06	-0.01	3.33	0.95	4.70	1.44	2.61	7.12	1.00
$NDE_g(2 \text{ vs } 1, S_{N_{ij}}(1))$	1.65	1.05	0.03	3.39	0.95	4.53	1.53	2.32	7.12	1.00
$NIE_g(2, S_{N_{ij}}(2) \text{ vs } S_{N_{ij}}(1))$	0.00	0.21	-0.36	0.35	0.51	0.20	0.48	-0.59	0.96	0.66
$NIE_g(1, S_{N_{ij}}(2) \text{ vs } S_{N_{ij}}(1))$	-0.01	0.20	-0.32	0.33	0.46	0.02	0.25	-0.28	0.46	0.45
$PCE_g(3 \text{ vs } 2)$	-0.03	0.28	-0.50	0.43	0.45	-2.39	1.29	-4.56	-0.46	0.02
$NDE_g(3 \text{ vs } 2, S_{N_{ij}}(3))$	0.00	0.00	0.00	0.00	0.00	-2.03	1.35	-4.35	0.02	0.05
$NDE_g(3 \text{ vs } 2, S_{N_{ij}}(2))$	0.00	0.00	0.00	0.00	0.00	-2.32	1.49	-4.74	-0.04	0.05
$NIE_g(3, S_{N_{ij}}(3) \text{ vs } S_{N_{ij}}(2))$	-0.03	0.28	-0.50	0.43	0.45	-0.07	0.60	-1.13	0.78	0.50
$NIE_g(2, S_{N_{ij}}(3) \text{ vs } S_{N_{ij}}(2))$	-0.03	0.28	-0.50	0.43	0.45	-0.36	0.60	-1.32	0.67	0.27
$PCE_g(3 \text{ vs } 1)$	1.62	1.03	0.03	3.31	0.95	2.34	1.04	0.69	3.99	0.99
$NDE_g(3 \text{ vs } 1, S_{N_{ij}}(3))$	1.64	1.11	-0.17	3.35	0.94	2.30	1.13	0.41	3.98	0.97
$NDE_g(3 \text{ vs } 1, S_{N_{ij}}(1))$	1.65	1.05	0.03	3.39	0.95	2.40	1.33	0.43	4.58	0.98
$NIE_g(3, S_{N_{ij}}(3) \text{ vs } S_{N_{ij}}(1))$	-0.03	0.30	-0.53	0.45	0.46	-0.07	0.62	-1.14	0.81	0.50
$NIE_g(1, S_{N_{ij}}(3) \text{ vs } S_{N_{ij}}(1))$	-0.02	0.49	-0.63	0.81	0.40	0.03	0.41	-0.43	0.73	0.41

Let us first focus on the contrast between Presentation and Flyer. We find evidence that principal causal effects of Presentation versus Flyer for Never Takers and Reward Compliers are positive. The posterior mean of the PCE for Never Takers is 1.65 (SD= 1.04) and the 90% credible interval only covers positive values; the posterior probability that this effects is positive is very high (96%). The posterior mean of the PCE for Reward Compliers suggests an increase of 4.73 museum visits due to presentation versus flyer (SD= 1.42 and 90%CI = (2.66, 7.14)) and this principal causal effect takes on positive values with probability one. Both results are consistent with those in Forastiere, Lattarulo, et al. (2021). Based on these

results, Presentation appears as an effective encouragement towards attending museums in the future for Reward Compliers and, to a smaller extent, for Never Takers.

As argued in Section 3.3, overall PCEs may originate from both a direct and an indirect/spillover causal pathway. Disentangling these pathways is our main contribution, moving a step forward with respect to the analyses in Forastiere, Lattarulo, et al. (2021) and enriching the literature on causal inference with interference. When comparing the Presentation and the Flyer, the direct pathway, whose strength is evaluated by principal NDEs, channels only encouragement effects for both Never Takers and Reward Compliers. Instead, the indirect pathway, whose strength is evaluated by principal NIEs, may only channel spillovers. Principal natural direct effects of Presentation versus Flyer can be evaluated setting the proportion of friends visiting Palazzo Vecchio to the value that it would have taken either under the Flyer encouragement, $S_{\mathcal{N}_{ij}}(1)$, or under the Presentation encouragement, $S_{\mathcal{N}_{ij}}(2)$. Similarly, principal natural indirect effects measure spillover effects on future museum attendance arising from a change in the proportion of friends visiting Palazzo Vecchio from $S_{\mathcal{N}_{ij}}(1)$, the value that $S_{\mathcal{N}_{ij}}$ would have taken under the Flyer encouragement ($z = 1$), to $S_{\mathcal{N}_{ij}}(2)$, the value that $S_{\mathcal{N}_{ij}}$ would have taken under the Presentation encouragement ($z = 2$), by fixing the individual encouragement to either the Flyer level ($z = 1$) or the Presentation level ($z = 2$). It is worthwhile to highlight that $S_{\mathcal{N}_{ij}}(1)$ is equal to the proportion of Always Taker friends, because only Always Takers visit Palazzo Vecchio under the Flyer encouragement; analogously, $S_{\mathcal{N}_{ij}}(2)$ is equal to the proportion of friends that are either Always Takers or Presentation Compliers, because only Always Takers and Presentation Compliers visit Palazzo Vecchio under the Presentation encouragement. Therefore, principal natural indirect effects quantify the spillover on future museum attendance that would come from friends ready to visit Palazzo Vecchio under the Presentation but not under the Flyer, that is, from being exposed, in addition to Always Takers, to friends that react as Presentation Compliers. Since the proportions of Always Takers and Presentation Compliers are small, $S_{\mathcal{N}_{ij}}(1)$ and $S_{\mathcal{N}_{ij}}(2)$ describe a friendship environment where a student is hardly

exposed to friends that might send spillovers.

The posterior means of $NDE_g(2 \text{ vs } 1, S_{\mathcal{N}_{ij}}(2))$ and $NDE_g(2 \text{ vs } 1, S_{\mathcal{N}_{ij}}(1))$ are 1.66 and 1.65 for Never Takers ($g = 000$) and 4.70 and 4.53 for Reward Compliers ($g = 001$), and the corresponding posterior probabilities that these effects are positive are higher than 0.95. We interpret these results as evidence that for Never Takers and Reward Compliers there exist pure encouragement effects of Presentation versus Flyer. The posterior means of $NIE_g(1, S_{\mathcal{N}_{ij}}(2) \text{ vs } S_{\mathcal{N}_{ij}}(1))$ and $NIE_g(2, S_{\mathcal{N}_{ij}}(2) \text{ vs } S_{\mathcal{N}_{ij}}(1))$ are very small for both Never Takers and Reward Compliers, and the corresponding 90% posterior credible intervals cover zero, therefore there is no evidence that there exist spillover effects of Flyer versus Presentation for these type of subjects. Indeed, as we can see in Table 8, the size of the estimated principal natural direct effects for Never Takers and Reward Compliers is similar to the size of the overall principal causal effects. Therefore, causal effects for Never Takers and Reward Compliers are almost entirely due to the encouragement.

Let us now focus on the contrast between Reward ($z = 3$) and Presentation ($z = 2$). For Never Takers, principal natural direct effects of Presentation versus Reward are zero by assumption (see the discussion in Sections 3.3 and 3.4), therefore principal causal effects of $z = 3$ versus $z = 2$ for this type of subjects are interpretable as spillover effects. Consistently with the insights of Forastiere, Lattarulo, et al. (2021), we find no evidence that there exist statistically significant spillover effects of Presentation versus Reward for Never Takers: the posterior means of $NIE_{000}(2, S_{\mathcal{N}_{ij}}(3) \text{ vs } S_{\mathcal{N}_{ij}}(2))$ and $NIE_{000}(3, S_{\mathcal{N}_{ij}}(3) \text{ vs } S_{\mathcal{N}_{ij}}(2))$ are very close to zero and the corresponding 90% posterior credible interval are spread around zero. For Reward Compliers, the posterior probabilities that overall principal causal effects of Presentation versus Reward are negative are close to 1, in line with the results previously presented in Forastiere, Lattarulo, et al. (2021). We interpret these overall principal causal effects for Reward Compliers as a mixture of experience and spillover effects, which we propose to disentangle using principal natural direct and indirect effects. Specifically, we attribute principal natural direct effects of Presentation versus Reward for Reward Compli-

ers solely to the experience (the visit at Palazzo Vecchio), because they visit Palazzo Vecchio only under Reward, and principal natural indirect effects to spillovers. Summary statistics of the posterior distributions of these effects shown in Table 8 suggest that experience effects account for almost all of the negative overall principal causal effect. Thus, Reward Compliers do not seem to have enjoyed their experience at Palazzo Vecchio very much, which may have even discouraged further museum visits in the follow-up period. If the effect of the Presentation itself is to enhance Reward Compliers’ museum attendance, such effect is partially offset by the experience at Palazzo Vecchio. Spillovers are, instead, negligible. This conclusion on spillovers is particularly relevant because, contrary to the Presentation/Flyer contrast, the Reward is able to increase the proportion of friends visiting Palazzo Vecchio, by attracting Reward Compliers, who are about 1/3 of students and may act as senders of spillovers.

Finally, let us focus on the contrast between Reward and Flyer. Overall principal causal effects of Reward versus Flyer amounts to the sum of the overall principal casual effects of Presentation versus Flyer and of Reward versus Presentation. We find evidence that for both Never Takers and Reward Compliers, these principal causal effects are positive and mainly due to direct effects: the posterior distributions of principal natural direct effects closely overlap the posterior distributions of overall principal causal effects, and posterior distributions of overall principal natural indirect effects are concentrated on values close to zero and evenly spread around zero.


5.3 Principal Controlled and Spillover Effects

Principal controlled direct effects measure the causal effects of the encouragement in situations where the proportion of friends performing the visit to Palazzo Vecchio is arbitrarily set by the researcher to a fixed value s^* for all units. In principle, any value $s^* \in [0, 1]$ could be chosen. To avoid excessive extrapolation, we focus on a range of proportions that are not unreasonably distant from those actually supported by the data. Figure 4 reports the CDEs

evaluated at six different levels of s^* .

Figure 4

Estimated CDEs



For Never Takers, it is worth showing only the CDE of the Presentation relative to the Flyer, $CDE_{000}(2 \text{ vs } 1, s^*)$. As shown in Table 5, the CDE of the Reward relative to the Presentation for Never Takers, $CDE_{000}(3 \text{ vs } 2, s^*)$, is assumed to be zero because the Reward applies only to students who actually visit Palazzo Vecchio, and disappears in the follow-up period. As a consequence, $CDE_{000}(3 \text{ vs } 1, s^*) = CDE_{000}(2 \text{ vs } 1, s^*)$. The posterior means of the CDEs of the Presentation relative to the Flyer for Never Takers show some evidence that there exists a positive direct effect of Presentation versus Flyer, which is approximately constant, equal to 1.65 additional museum visits, irrespective of the selected proportion of friends visiting Palazzo Vecchio. The 90% posterior predictive intervals cover zero, even if zero is very close to the lower bounds of the intervals, especially for proportion of friends visiting Palazzo Vecchio lower than (or equal to) 0.3.

For Reward Compliers, it makes sense to consider CDEs relative to any encouragement contrast (see previous discussion and Table 5). $CDE_{001}(2 \text{ vs } 1, s^*)$ can be interpreted as a pure encouragement effect, $CDE_{001}(3 \text{ vs } 2, s^*)$ as the effect of the visit experience at Palazzo Vecchio, and $CDE_{001}(3 \text{ vs } 1, s^*)$ as the sum of the previous two quantities. The CDEs of

the Presentation relative to the Flyer for Reward Compliers are to raise the number of subsequent visits by more than 4 points, regardless of the fixed proportion of friends visiting Palazzo Vecchio. Therefore, the presentation has the effect of raising the intrinsic motivation of Reward Compliers to visit museums in the follow-up period. The posterior mean of the CDEs of Reward versus Presentation for Reward Compliers is negative but statistically negligible; the 90% posterior credible intervals cover zero and are quite large. Therefore, there is no evidence that the Palazzo Vecchio experience may prompt future museum visits for Reward Compliers.

Finally, Figure 5 shows the means and the 90% credible intervals of the posterior distributions of CSEs, originating from some discretionary change in the proportion of friends visiting Palazzo Vecchio while the encouragement is kept at a fixed level. In particular, the CSEs are evaluated at some alternative variations of such proportion from/to an arbitrary level that resembles the average proportion of friends visiting Palazzo Vecchio under the encouragement considered. This allows to evaluate how spillovers change when the proportion of friends visiting Palazzo Vecchio is higher or lower with respect to the latter.


The spillover generated by any of these changes on the future museum attendance of both Never Takers and Reward Compliers are small and statistically negligible: the posterior means are close to zero and the 90% posterior credible intervals cover zero and are approximately symmetric around zero. Therefore, no matter how much an individual is exposed to friends visiting Palazzo Vecchio, spillovers are hardly relevant in the decision to attend museums in the follow-up period.

6 Concluding Remarks

In this article we show how information on social networks can be used to investigate spillover effects within cluster encouragement designs with individual non-compliance. We propose an original methodological interplay between principal stratification and mediation analy-

Figure 5

Estimated CSEs for Never Takers and Reward Compliers



Note. The blue dashed line indicates zero.

sis, where a synthesis of the network information is used as a mediator which may generate spillover effects. We formally define principal natural direct and indirect effects and principal controlled direct and spillover effects, and explain how these quantities can be interpreted to get information on different causal pathways, including spillover. We revisit results from a small field experiment, based on a clustered encouragement design, conducted in Florence (Italy) to study how appropriate personal incentives may lead students to visit museums in their free time. Possible causal pathways can either originate from such incentives directly affecting the student's outcome, or follow an indirect trail where the student's outcome is

finally affected by the spillovers received from other students involved in the experiment. We use a Bayesian approach to inference under latent sequential ignorability assumptions. Previous causal studies of these experimental data (Lattarulo et al., 2017; Forastiere, Lattarulo, et al., 2021) have not exploited the available information on friendship networks within each cluster of students, which we use here to disentangle and estimate spillover effects, and have thus settled for causal effects where spillovers are usually blended with other sources of change in the outcome. Since our application is based on a small field experiment, the conclusions we reach would perhaps require corroboration in similar studies conducted on a larger scale. However, the results of our analysis suggest that spillovers from friends did not lead to significant changes in the students’ museum attendance, the Presentation encouragement was effective in motivating students to devote some free time to museum attendance, the visit experience that students were offered at Palazzo Vecchio was not effective, at least for the subpopulation of Reward Compliers.

Acknowledgements

The authors thank schools, teachers and students who participated in the study, as well as the educational experts of the Mus.e Association for having performed the classroom presentations. The authors would also like to thank Laura Forastiere for her valuable comments.

References

- Angrist, J. D., Imbens, G. W., & Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434), 444–455.
- Cellini, R., & Cuccia, T. (2018). How free admittance affects charged visits to museums: An analysis of the italian case. *Oxford Economic Papers*, 70(3), 680–698.
- Cochran, W. (1963). *Sampling techniques*. 2nd edition. John Wiley & Sons.

- Cox, D. (1958). *Planning of experiments*. Wiley.
- Del Prete, D., Forastiere, L., & Leone Sciabolazza, V. (2019). Causal inference on networks under continuous treatment interference: an application to trade distortions in agricultural markets. *Available at SSRN 3363173*.
- Forastiere, L., Airoidi, E. M., & Mealli, F. (2021). Identification and estimation of treatment and interference effects in observational studies on networks. *Journal of the American Statistical Association*, 116(534), 901–918.
- Forastiere, L., Lattarulo, P., Mariani, M., Mealli, F., & Razzolini, L. (2021). Exploring encouragement, treatment, and spillover effects using principal stratification, with application to a field experiment on teens’ museum attendance. *Journal of Business & Economic Statistics*, 39(1), 244–258.
- Forastiere, L., Mealli, F., & VanderWeele, T. J. (2016). Identification and estimation of causal mechanisms in clustered encouragement designs: Disentangling bed nets using bayesian principal stratification. *Journal of the American Statistical Association*, 111(514), 510–525.
- Forastiere, L., Mealli, F., Wu, A., & Airoidi, E. (2018). Estimating causal effects under interference using bayesian generalized propensity scores. *arXiv preprint arXiv:1807.11038*.
- Frangakis, C. E., & Rubin, D. B. (2002). Principal stratification in causal inference. *Biometrics*, 58(1), 21–29.
- Frangakis, C. E., Rubin, D. B., & Zhou, X.-H. (2002). Clustered encouragement designs with individual noncompliance: Bayesian inference with randomization, and application to advance directive forms. *Biostatistics*, 3(2), 147–164.
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). *Bayesian analysis*, 1(3), 515–534.

- Goldsmith-Pinkham, P., & Imbens, G. W. (2013). Social networks and the identification of peer effects. *Journal of Business & Economic Statistics*, 31(3), 253–264.
- Gustafson, P. (2010). Bayesian inference for partially identified models. *The international journal of biostatistics*, 6(2).
- Hirano, K., Imbens, G. W., Rubin, D. B., & Zhou, X.-H. (2000). Assessing the effect of an influenza vaccine in an encouragement design. *Biostatistics*, 1(1), 69–88.
- Hong, G., & Raudenbush, S. W. (2006). Evaluating kindergarten retention policy: A case study of causal inference for multilevel observational data. *Journal of the American Statistical Association*, 101(475), 901–910.
- Hudgens, M. G., & Halloran, M. E. (2008). Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482), 832–842.
- Hughes, K., & Moscardo, G. (2019). For me or not for me? exploring young adults’ museum representations. *Leisure Sciences*, 41(6), 516–534.
- Imbens, G. W., & Rubin, D. B. (1997). Bayesian inference for causal effects in randomized experiments with noncompliance. *The annals of statistics*, 305–327.
- Kang, H., & Imbens, G. (2016). Peer encouragement designs in causal inference with partial interference and identification of local average network effects. *arXiv preprint arXiv:1609.04464*.
- Kisida, B., Greene, J. P., & Bowen, D. H. (2014). Creating cultural consumers: The dynamics of cultural capital acquisition. *Sociology of Education*, 87(4), 281–295.
- Lattarulo, P., Mariani, M., & Razzolini, L. (2017). Nudging museums attendance: a field experiment with high school teens. *Journal of Cultural Economics*, 41(3), 259–277.
- Manna, R., & Palumbo, R. (2018). What makes a museum attractive to young people? evidence from Italy. *International Journal of Tourism Research*, 20(4), 508–517.

- Mattei, A., & Mealli, F. (2007). Application of the principal stratification approach to the faenza randomized experiment on breast self-examination. *Biometrics*, 63(2), 437–446.
- Morris, S. S., Flores, R., Olinto, P., & Medina, J. M. (2004). Monetary incentives in primary health care and effects on use and coverage of preventive health care interventions in rural honduras: cluster randomised trial. *The Lancet*, 364(9450), 2030–2037.
- Ogburn, E. L., Sofrygin, O., Diaz, I., & Van Der Laan, M. J. (2017). Causal inference for social network data. *arXiv preprint arXiv:1705.08527*.
- Pearl, J. (2001). Direct and indirect effects. In *Proceedings of the seventeenth conference on uncertainty in artificial intelligence* (pp. 411–420).
- Robins, J. M., & Greenland, S. (1992). Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 143–155.
- Rubin, D. B. (1980). Randomization analysis of experimental data: The fisher randomization test comment. *Journal of the American statistical association*, 75(371), 591–593.
- Sobel, M. E. (2006). What do randomized studies of housing mobility demonstrate? causal inference in the face of interference. *Journal of the American Statistical Association*, 101(476), 1398–1407.
- Stan Development Team. (2014). Stan: A c++ library for probability and sampling. *Online*: <http://mc-stan.org>.
- Ten Have, T. R., & Joffe, M. M. (2012). A review of causal estimation of effects in mediation analyses. *Statistical Methods in Medical Research*, 21(1), 77–107.
- VanderWeele, T. J., & An, W. (2013). Social networks and causal inference. *Handbook of causal analysis for social research*, 353–374.
- VanderWeele, T. J., & Vansteelandt, S. (2009). Conceptual issues concerning mediation, interventions and composition. *Statistics and its Interface*, 2(4), 457–468.