



Peer effects, self-selection and dishonesty

Liza Charroin, Bernard Fortin, Marie Claire Villeval

► To cite this version:

Liza Charroin, Bernard Fortin, Marie Claire Villeval. Peer effects, self-selection and dishonesty. Journal of Economic Behavior and Organization, 2022, 200, pp.618-637. 10.1016/j.jebo.2022.06.024 . hal-03712450

HAL Id: hal-03712450

<https://hal.science/hal-03712450>

Submitted on 3 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Peer effects, self-selection and dishonesty[☆]

Liza Charroin^a, Bernard Fortin^b, Marie Claire Villeval^{c,d,*}

^a University Paris 1 Panthéon Sorbonne, Centre d'Economie de la Sorbonne, Paris, France

^b Economics Department, Laval University Québec, CIRPÉE, and CIRANO, Canada

^c Univ Lyon, CNRS, GATE UMR 5824, 93 Chemin des Mouilles, Ecully, F-69130, France

^d IZA, Bonn, Germany

ARTICLE INFO

Article history:

Received 9 February 2022

Revised 20 June 2022

Accepted 22 June 2022

JEL classification:

C92

D83

D85

D91

Keywords:

Peer effects

Self-selection

Homophily

Dishonesty

Experiment

ABSTRACT

If individuals tend to behave like their reference group, is it because of peer effects, self-selection, or both? Using a peer effect model allowing for conformity and link formation, we designed a real-effort laboratory experiment in which individuals could misreport their performance and select their peers. Our results reveal both a preference for conformity and homophilous link formation, but only among individuals cheating in isolation. This suggests that such link formation was not motivated by a taste for similarity but by acquiring self-serving information. Importantly, we reject the presence of a self-selection bias in the peer effect estimates by showing that the size of peer effects is similar when identical peers were randomly assigned and when individuals selected them.

© 2022 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY license
(<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

Most parents care a lot about their children's friendship networks because they fear the negative influence that some peers may exert on their children's behavior. Indeed, there is evidence that unethical behavior disseminates within social networks. For example, some corporate cultures favor weak ethical norms (Egan et al., 2019), certain social interactions encourage criminal acts (Glaeser et al., 1996), and having unethical peers in an academic context increases the individuals' temptation to cheat (Carrell et al., 2008). In the ethical domain, as in other contexts, peer effects may result from a preference for conformity – namely, the tendency of people to align behavior with the behavior of their peers –, because deviating

[☆] We thank Vincent Boucher, Yann Bramoullé, Yan Chen, Alain Cohn, Russell Davidson, Claude Fluet, Steeve Marchand, Simeon Schudy and Bruce Shearer for helpful discussions. We also thank seminar participants at Aix Marseille School of Economics, Beta at Université de Lorraine, CREST in Paris, the Erasmus Institute in Rotterdam, ICES at George Mason University, Laval University, Université du Québec Montréal, the University of Michigan SBEE seminar, Purdue University, Shandong University, the Universities of Glasgow, Munich and Heidelberg, and Virginia Tech University for helpful comments. Aristide Houndetoungan provided exceptional research assistance. We are also grateful to Quentin Thevenet who programmed the experiment. Support for this research was provided by the French National Research Agency (FELIS, ANR-14-CE28-0010-01), the CORTEX Laboratory of Excellence (ANR-11-LABX-0042) of Université de Lyon, within the program Investissements d'Avenir (ANR-11-IDEX-007), and IDEXLYON from Université de Lyon (INDEPTH-IDEX/SBP/2018/03) within the Programme Investissements d'Avenir (ANR-16-IDEX-0005) operated by the French National Research Agency.

* Corresponding author.

E-mail addresses: liza.charroin@univ-paris1.fr (L. Charroin), bernard.fortin@ecn.ulaval.ca (B. Fortin), villeval@gate.cnrs.fr (M.C. Villeval).

from others' behavior generates a disutility, while conformity brings advantages (Henrich and Boyd, 1998). However, similar behavior in a network may also result from homophily, that is, the individuals' tendency to create more links with others who have similar socio-demographic characteristics, such as gender or ethnicity, or similar aspirations, values or moral preferences (McPherson et al., 2001; Benhabib et al., 2010; Bramoullé et al., 2012; Baccara and Yariv, 2013; Acemoglu et al., 2021).

When analyzing behavior in a network, one may interpret the correlation of behavior across individuals as peer effects, while it stems from the choice of similar peers. Thus, the presence of a *self-selection bias* may overestimate peer effects. Such a bias is not present when the network is stochastic but exogenous. However, the formation of networks is usually endogenous with observational data. This makes the identification and estimation of causal peer effects particularly challenging. In that case, consistent estimators of peer effects require sophisticated econometric techniques that are still being developed (e.g., de Paula, 2017; Graham, 2019). Unfortunately, most papers ignore the self-selection bias and may thus provide inconsistent peer estimators. Furthermore, the recent studies that attempt to account for this bias have reached mixed conclusions. While some of them suggest that there is no severe endogeneity bias (e.g., Goldsmith-Pinkham and Imbens, 2013; Boucher, 2016; Badev, 2021), others reach the opposite conclusion (e.g., see Carrell et al., 2013; Hsieh et al., 2020a).

Our paper addresses these crucial issues thanks to a laboratory experiment in which the matching process is perfectly observable and properly manipulated. In particular, we designed an experiment where the network is stochastic but exogenous, that is, randomized. The experimental data using this exogenous network is used to easily estimate a model with no self-selection bias, and to obtain consistent causal peer effect estimators. In another experiment where the network is stochastic but endogenous, individuals select their own peers.

The main objective of our research is to test whether the selection of peers biases the measure of peer effects in an ethical dilemma context. We chose this context because the self-selection of peers may have significant consequences in this domain. Indeed, through the selection of peers, individuals may also choose the social norm that they would like to comply with, which may help them justify their own (mis)behavior. We first identify pure conformity effects on lying behavior in a setting where various types of peers (characterized by a signal on their moral type) is *exogenously* assigned to individuals. We investigate whether conformity is affected by the type of peers, that is, whether there is an observable heterogeneity in peer effects (e.g., Xiang et al., 2010; Aral and Walker, 2014; Patacchini et al., 2017). Given the nature of the outcome, peer effects are limited to conformity, with no complementarity or social learning in the lying production function. Second, we evaluate homophilous link formation by letting individuals select their peers. What we call "homophilous link formation" in this paper is a preference for being matched with peers who made a similar initial choice that sends a signal on their moral values.¹ A homophilous link formation in our context may result from a pure taste for similarity or from the strategic selection of a given channel of information. Finally, when individuals can choose their peers, we test if the tendency to select similar peers (if any) affects the estimate of conformity by comparison with the setting in which peers are exogenously assigned. Note that the difference between conformity estimates when networks are exogenous and endogenous can stem not only from a homophilous link formation but also from the presence of a *link strength effect* in the latter case. A link strength effect occurs if individuals weigh more the examples of peers because they have selected them. Overall, such homophilous link formation and link strength effect may artificially inflate the estimated conformity effects when peers are endogenously selected.

By answering these questions, we contribute to the peer effects literature on unethical behavior, using observational data (e.g., Patacchini and Zenou, 2009; Damm and Dustmann, 2014; Ajzenman, 2021), or experimental data (e.g., Fortin et al., 2007; Keizer et al., 2008; Kroher and Wolbring, 2015; Kocher et al., 2018; Bott et al., 2020; Drago et al., 2020). Compared to this empirical literature, we innovate by measuring the extent to which the endogenous network formation affects conformity effects. We also innovate by exploring how selective matching in endogenous networks allows individuals to acquire information selectively in the ethical domain. Surprisingly, very few experimental studies on cheating have explored the role of matching (Gross et al., 2018; Akin, 2019; Leib and Schweitzer, 2020), while sorting opportunities have been found to affect behavior in other domains, such as social preferences (Lazear et al., 2012), public goods (Page et al., 2005), and education (Kiessling et al., 2019). In contrast to Gross et al. (2018), in our study there is neither complementarity in payoffs nor social learning. Thus, we can study pure conformity effects, when the formation of networks is exogenous or endogenous, to assess the influence of self-selection on the size of peer effects.

In our experiment, before performing a simple task repeatedly for a piece rate, participants have to choose, only once, between two modes of calculation of their earnings performance.² Under the *Automatic mode*, the program directly computes the participants' actual performance. Under the *Manual mode*, the participants have to compute themselves and self-report their performance, which provides them some discretion to increase their earnings by cheating. This choice opportunity is a novel aspect of our design (for an exception, see Konrad et al., 2021) because previous studies only observed how people succumb to temptation when opportunities are given to them. In our design, the mode chosen by an individual provides an (imperfect) signal about the morality of this individual. Indeed, the Manual mode requires additional costly effort, and

¹ Homophily is defined in the literature as the tendency for individuals to bond with similar others. As noted earlier, the similarity is not necessarily defined in terms of socio-demographic characteristics but also in terms of characteristics such as attitudes and aspirations.

² Outside the laboratory, individuals cannot usually decide how their performance is evaluated. But they can self-select into occupations or sectors that give more or less unethical opportunities for personal enrichment.

it should be chosen mainly for its cheating opportunity. Thus, it should be considerably more attractive for individuals who have weaker moral values than for more honest individuals.

In part 1 of all treatments, individuals perform the task in isolation. The content of part 2 differs across three between-subjects treatments: while in the Baseline, individuals perform again the task in isolation, in the Exogenous (EXO) and the Endogenous treatments (ENDO) they receive social information about peers' performance from the Baseline before executing the task. The Baseline is run first to collect data for the two subsequent treatments. At the beginning of part 2 in EXO and ENDO, participants are matched with two peers from the Baseline: either peers who both chose the Automatic mode or peers who both chose the Manual mode. At the beginning of each period, individuals matched with peers who chose the Automatic mode are informed of these peers' average *actual* performance in the same period; individuals matched with peers who chose the Manual mode are informed of their average *reported* performance.

Creating unidirectional networks with a one-way flow of information removes any simultaneity (or reflection, using Manski (1993)'s terminology) issue. This is an advantage of a laboratory setting, which also minimizes endogeneity associated with measurement errors and partial sampling of the network (e.g., de Paula, 2017). In EXO, each participant is randomly matched with peers, allowing us to estimate a pure conformity effect with no bias, and test whether the effect is heterogeneous depending on the type of peers. In ENDO, participants have to choose their two peers. They have to select either a pair of peers who chose the Manual mode or a pair of peers who chose the Automatic mode and are also informed about the average performance of each pair in part 1. The mode provides a signal of the peers' moral values, which allows us to test the presence of homophilous link formation. Then, by comparing the conformity effects in EXO and ENDO using a joint test, we can identify the presence of a positive selection bias in the peer effects or of a link strength effect.

To estimate conformity effects, we develop a quadratic social interactions model inspired by Ballester et al. (2006), Calvó-Armengol et al. (2009) and Boucher and Fortin (2016), where the individual's utility depends on net payoffs, the moral cost of cheating, and average peers' performance and type. In ENDO, the model is augmented to account for the endogenous choice of peers, which requests introducing a selection equation providing the likelihood of selecting peers who chose the same mode as the individual.

The results of EXO provide evidence of peer effects on lying. Individuals did not lie maximally in isolation and the size of lies of those who chose the Manual mode increased with their peers' reported performance. This was observed only for individuals with a weak moral type (as revealed by their mode choice and lying behavior in isolation) and only when they were matched with peers who chose the same mode. Participants who were matched with Manual mode peers conformed more to the reported performance of these peers as if it reduced their moral cost of lying. By contrast, those who chose the Manual mode but did not lie in isolation did not lie more after observing the reported performance of peers who potentially lied. They treated this information as irrelevant.

The most original results are the following. First, there is evidence of homophilous link formation in ENDO, but again, only among individuals with a weaker moral type. When matching was endogenous, liars were more likely to create a link with peers who also chose the Manual mode than peers who chose the other mode. By contrast, individuals who chose the Manual mode and did not lie in isolation and those who chose the Automatic mode did not express any preference in matching. This difference is informative on the mechanism driving such link formation. If the mechanism was a pure taste for similarity, we should have observed them in both liars and honest participants. The fact that only liars created homophilous links suggests that they were motivated by the willingness to form a reference group whose behavior represents a lenient descriptive social norm. Second, comparing peer effects in EXO and ENDO reveals no significant difference in the size of peer effects when the network formation was exogenous and when it was endogenous. This test jointly rejects the presence of a self-selection bias in the estimation of peer effects on lying and the existence of a link strength effect. Therefore, as long as one considers settings without strategic complementarity, our results are consistent with studies that conclude that homophilous link formation is not the source of a severe bias on peer estimators.

The remainder of this paper is organized as follows. Section 2 reviews the related literature. Section 3 describes our experimental design. Section 4 develops our theoretical model. Section 5 presents our results. Finally, Section 6 discusses and concludes.

2. Related literature

We contribute to three strands of the literature: the identification of peer effects when networks are formed endogenously, the identification of homophilous link formation, and the role of peers in the dissemination of unethical behavior.

First, our study complements the literature on peer effects based on observational data, which has notably revealed the role of conformity on crime (Glaeser et al., 1996; Patacchini and Zenou, 2009; Damm and Dustmann, 2014; Daz and Patacchini, 2020). However, most studies do not address the issue of the endogeneity of networks, which may bias the identification of peer effects. Various econometric techniques have been recently developed to address this endogeneity issue (see de Paula, 2017, for a recent survey) but they have not been applied to unethical behavior. That individuals self-select into a network to optimize their utility function has been considered recently. Boucher (2016) proposed a social interaction model with conformity, estimated using an Add Health sample of students who simultaneously chose their participation in extracurricular activities and peers to connect to. He found that network endogeneity has little impact on peer effects. Testing a model in which individuals both choose their peers and the decision to smoke tobacco, Badev (2021) also found that accounting for the endogeneity of network formation has little impact on the peer effect




Fig. 1. Timeline of the experiment.

parameter. Hsieh et al. (2020b) showed that students create more links with good students to increase their chance of academic success. These studies reveal a weak effect of the endogeneity of network formation on behavior, but the results are sensitive to the econometric approaches used.

Exploring peer effects on ethics with observational data is particularly challenging when one is willing to account for link formation. At the cost of increased artificiality, using a lab experiment allows us to test whether the formation of links influences the measure of peer effects with perfect observability of networks and behavior, and with econometric methods that are easier to implement.

Second, our study contributes to the literature on homophily (Currarini et al., 2009; Benhabib et al., 2010; Golub and Jackson, 2012). Since the study of McPherson et al. (2001) showing that similarity in terms of characteristics, values, or beliefs generates connections, most experimental results on homophily are from studies of group identity (for a survey, see Li, 2020). For example, in Currarini and Mengel (2016), individuals preferably match with in-groups and homophilous individuals are less reciprocal toward out-groups than in-groups. There is also evidence of assortative matching in the field of criminality (e.g., Haynie, 2001; Knecht et al., 2010; Gavrilova, 2019). Flashman and Gambetta (2014) showed that deviants are more likely to link with other deviants than with non-deviants because it creates mutual obligations. Here, we consider a context of unethical decision-making without productive complementarity and with matching that is based on a signal of moral values. Moreover, we aim to analyze whether the selection of peers influences the size of causal peer effects and biases their estimation.

Also, we renew the experimental literature on peer effects on cheating. While some studies found that observing peers' behavior increases dishonesty,³ others found little effects.⁴ We provide novel evidence in a different setting. Crucially, in almost all past studies, peers were exogenously assigned, whereas outside the laboratory individuals can usually choose their peers. As one rare exception, Gross et al. (2018) allowed individuals to change partners in a repeated dyadic task and found that people rematched opportunistically. By contrast, we compare lying in isolation and in the presence of exogenous and endogenous social interactions, allowing us to evaluate self-serving peer selection. We differ from Leib and Schweitzer (2020) who endogenized the access to free or costly social information and focused on the duration of the search, and from Pascual-Ezama et al. (2015) who invited real friends and demonstrated that peer effects depend on the number of peers cheating. In our study, all participants receive the same quantity of information, so that we are able to isolate the pure effect of the endogeneity of matching. Finally, we differ from Akin (2019) who let participants choose to be informed either on the average score or the maximum score of peers, and showed that sorting exacerbates dishonesty. By contrast, we can identify who lied and who did not, which allows us to measure assortative matching.

3. Experimental design

We start by describing the task performed by the participants in the two parts of the experiment. Next, we introduce the three treatments that allow us to measure peer effects, homophilous link formation, and the existence of a self-selection bias or a link strength effect in the estimation of causal peer effects. The two parts were instructed separately. Fig. 1 displays the timeline of the experiment.

³ See Keizer et al. (2008); Gino et al. (2009); Fosgaard et al. (2013). Information on the dishonesty of others increases lying: see Innes and Mitra (2013) in a sender-receiver game, Baker and Mechtel (2019) and Lauer and Untertrifaller (2019) in real-effort tasks. Similarity increases cheating in the presence of complementarities (Irlenbusch et al., 2020). But regardless of payoff commonality, individuals lie more in groups because they can communicate (Kocher et al., 2018). Drago et al. (2020) showed that informing targeted individuals on the risk of sanctions can generate spillover effects on untreated neighbors.

⁴ See Fortin et al. (2007); Rauhut (2013); Kroher and Wolbring (2015); Abeler et al. (2019); bad examples are more contagious than good ones (Lefebvre et al., 2015; Dimant, 2019); and norm compliance may depend on the proximity with peers (Bicchieri et al., 2020).

3.1. The task

In all treatments, participants have to perform the Counting-0s task introduced by Falk et al. (2006). In each period, they have two minutes to count the number of 0s in grids that only contain 0s and 1s. This task does not require any prior knowledge. Its purpose is to capture differences in effort, not ability, thus eliminating social learning. Grids appear one by one on the screens. After an answer is validated, correct or not, a new grid is displayed. Individuals could see a maximum of 20 grids in each period (only two participants over 369 managed to solve 20 grids in two minutes; they did so only once). For a reason explained later, the grids have two sizes, each associated with a different piece rate: 5×5 grids paid 1 point, and 5×10 grids paid 2 points if solved correctly (see Figure A.1 and A.2). The sequence of grids was randomly generated before the experiment, and every participant received the same grids in the same order, such that actual performances are perfectly comparable across individuals.

We used this task instead of a random task such as the die task because even in trivial tasks that involve the provision of effort, we expected a majority of individuals to value their performance, which may make them more sensitive to social comparisons. At the end of each period, a feedback table is displayed on the participant's screen (see Figures A.3 and A.4). Each line corresponds to a grid with, in columns, its piece rate, the answer provided, and the correct answer. We define performance in a period as the sum of points earned in this period. Thus, the actual performance is obtained by summing the number of grids solved, multiplied by their corresponding piece rate. This task is repeated (with different grids) in each period of the two parts.

3.2. Part 1: Mode choice and task performance in isolation

Part 1 is similar across treatments. At the beginning of part 1, participants practice the task for two minutes, and they observe an example of the feedback table. Next, they have to choose between two modes of performance evaluation: the Automatic and the Manual modes.⁵ Choosing the *Automatic mode* means that their performance is computed and recorded by the program. Choosing the *Manual mode* means that they have to calculate and report themselves their performance on the computer. Based on the feedback table (see Figure A.4 in Appendix), they have to check whether they solve each grid and sum the number of points earned depending on the corresponding piece rates. This action requires additional effort but this mode also creates a cheating opportunity because participants can then misreport their performance. We used two grid sizes to increase the opportunity (cognitive) cost of choosing the Manual mode if an individual does not intend to cheat. Thus, participants have to not only check whether each answer is correct but also apply the piece rate corresponding to the grid size. If this mode is selected, we can measure the actual performance and the size of lies, provided by the difference between the actual and the reported performances.

As for the task, cheating does not require any sophistication, which should exclude that peer effects, if any, are driven by social learning on cheating. Participants had to figure out by themselves that over-reporting is feasible. We did not explicitly inform them that they could cheat to avoid introducing an experimenter demand effect. This is standard in experiments studying cheating behavior. For the same reason, we did not inform them that we could observe the size of their lies, but they could infer it from the feedback that indicates their actual performance for each grid.

After choosing the mode that applies throughout the part,⁶ participants perform the task in isolation during five periods of two minutes each. At the end of each period, depending on the mode, they have or not to report the number of points earned.

3.3. Part 2: Introduction of social interactions

At the beginning of part 2, participants learn that they have to perform the same task with new grids during five periods of two minutes each, under the same mode as in part 1. Part 2 differs from part 1 across three between-subjects treatments. In the Baseline, individuals perform again the task in isolation. In EXO, they are exogenously matched with two peers from the Baseline who selected either the Automatic or the Manual mode, and they are informed of the mode and the average performance of these peers at the beginning of each period. ENDO is similar, except that participants have to choose their two peers. Pairs provide a sufficient minimum number of peers allowing us to test linear-in-means models. We now describe each treatment.

In the Baseline, after a five-minute break after part 1 to mimic the timeline of the other treatments (participants had to wait in silence), participants perform again the task in isolation. These participants represent the pool of peers that are matched with participants in the other treatments. For that reason, the data of the Baseline were collected first. By matching participants in the Baseline with participants in the other treatments, we form unidirectional networks where information

⁵ In the instructions (see Appendix A1), the Automatic mode was called the Direct mode and the Manual mode was called the Indirect mode. The instructions used neutral language.

⁶ We acknowledge that since participants had to choose their mode at the beginning of part 1 once for all, some of them may have regretted their choice later on. Someone who chose the Manual mode remained free to lie or to report honestly later on, whereas the choice was irreversible for those who selected the Automatic mode but could be tempted to lie later on. Thus, we can identify the moral type of these individuals imperfectly, only through their single choice of mode.

on actual or reported performance flows one-way. This avoids any simultaneity issues in the estimation of peer effects. Participants were informed that their performance data (always kept anonymous) could be used in future sessions.

In EXO, at the beginning of part 2, participants are exogenously matched with peers from the Baseline according to the following procedure. Each one first receives information on the average performance in part 1 of two pairs of peers from the Baseline. One pair comprises peers who chose the Automatic mode (thus, the numbers indicate their average *actual* performance), and the other pair comprises peers who chose the Manual mode (thus, the numbers indicate their average *reported* performance). We provide these two averages to make it clearer to the participants that the Manual mode leads to a reported performance probably inflated by cheating.⁷ Importantly, participants were never informed on the percentage of peers who selected each mode. Next, each participant is randomly and equiprobably matched with one of these two pairs. These unidirectional networks remain fixed throughout part 2. Thus, there are four categories of individuals defined by the mode(s) chosen by the participants and by their peers.

At the beginning of each period, the participants are informed of the average performance of their two peers in the same period of part 2 (when seeing the same grids). It was important to provide new information in each period to avoid an anchoring effect and to take into account the possible dynamic of peers' performance over time. To be consistent with the properties of a linear-in-means model of peer effects, we provide them the average performance instead of the performance of each peer separately. If the participants are matched with peers who chose Automatic (respectively, Manual), it is common knowledge that they observe the average actual (respectively, reported) performance of the peers. If they are matched with peers who chose Manual, they are not informed of the existence and size of peers' lies.⁸ After observing the peers' performance, participants have to perform the task. Payment is based on the individual performance.

Similar to EXO, participants in ENDO perform the task in isolation in part 1, they are matched with two peers in part 2, and they are informed on the average performance of these two peers in the same period at the beginning of each of the five periods of part 2. In contrast with EXO, they have to choose between two pairs of peers at the beginning of part 2. They learn the average actual performance in part 1 of a pair of peers who selected the Automatic mode and the average reported performance of a pair of peers who chose the Manual mode. By selecting one pair, they either choose to observe the average actual performance in each new period of the peers that selected the Automatic mode or the average reported performance of the peers that selected the Manual mode.^{9,10}

3.4. Procedures

The experiment was conducted at GATE-Lab, Lyon, France. The 369 participants, mostly undergraduate students with majors in local engineering, business, and medicine, were recruited online using Hroot (Bock et al., 2014). We ran four sessions with the Baseline ($N = 72$) and then six sessions with EXO ($N = 143$) and six sessions with ENDO ($N = 154$) (see Table A.1 and Table A.2 on individual characteristics in Appendix). EXO and ENDO requested more participants because the protocol generates four categories of participants instead of two in the Baseline because of the matching procedure.

Upon arrival, participants were assigned to a terminal after drawing a computer tag. The instructions for each part were distributed at the beginning of each part and read-aloud (see Appendix A1). They were written in neutral language. Before starting the first part, participants had to fill out a comprehension questionnaire (see Appendix A2) and practised the task for two minutes. Participants had to figure out by themselves that over-reporting was feasible and we did not explicitly inform them that we were able to measure the size of their lies (but they could figure this out since we provided them with the actual performance feedback for each grid). Participants knew from the beginning that there would be several parts, but they did not know what would happen in part 2 before part 1 was completed. Instructions for part 2 were distributed only after part 1 was completed. At the end of the session, one period in each part was randomly drawn for payment. As pre-announced in the instructions to avoid the selection of a given pair of peers by pure curiosity, players in EXO or ENDO received feedback on the mean performance in each of the five periods of the pair of peers they were not matched with in part 2. Finally, after they filled out a sociodemographic questionnaire, participants were paid in private in a

⁷ Pairs can not be of mixed type because the information about average performance would be harder to interpret. In EXO, the observed average actual performance of pairs who chose Automatic in part 1 was 12.84 points (min = 6.2, max = 18.8) and the average reported performance of pairs who chose Manual was 21.26 points (min = 11.3, max = 38.9). In 90% of the cases, individuals observed an average reported performance of the pair that chose Manual greater than the average actual performance of the pair that chose Automatic. The mean difference between the two averages was 8.42 points.

⁸ We acknowledge that this induces some ambiguity on the intentions of peers. We chose to not inform players about the size of peers' lies to avoid an experimenter demand effect. Indeed, this would have focused the attention of players on lying opportunities, inducing another difference with the instructions for part 1. Moreover, outside of the lab, one can usually only suspect the existence of cheating, like on credence goods markets. It would have been also difficult to elicit the participants' beliefs on the level of lying of their peers without creating a demand effect.

⁹ In ENDO, the observed average actual performance in part 1 of pairs that chose Automatic was 12.91 points (min = 6.1, max = 18.5), and the average reported performance of pairs that chose Manual was 20.55 points (min = 12, max = 38.9). In 91% of the cases, individuals in ENDO observed an average reported performance of the pair that chose Manual greater than the average actual performance of the pair that chose Automatic. The mean difference between the two averages was 7.64 points. There is no significant difference with the statistics in the EXO treatment detailed in Footnote.⁷

¹⁰ To avoid that the choice was driven by curiosity about the reported performance of potential cheaters, participants were informed upfront that they would receive at the end of the session information about the mean performance in each period of the pair of peers that they did not select. Thus, if individuals choose peers who selected the Manual mode, we can exclude that it is driven by pure curiosity. For symmetry reasons, we also provided this information in EXO.

separate room. On average, sessions lasted approximately 75 minutes. Participants earned on average 16.36 Euros (standard deviation = 3.73), including a 5-Euro show-up fee.

4. Theoretical model

We start by presenting a formal model of the mode choice, performance, and lying decisions in isolation. Next, we generalize the model to account for exogenous and endogenous social interactions.

4.1. Mode choice and lying in isolation

The model is solved recursively. First, conditional on choosing the Automatic mode, we assume a quadratic utility function (see Ballester et al., 2006; Calvó-Armengol et al., 2009; Boucher and Fortin, 2016) given by

$$U_{it} = (\mathbf{x}'_i \boldsymbol{\beta} + \epsilon_{it}) p_{it} - \frac{p_{it}^2}{2}, \quad (1)$$

where p_{it} denotes the actual performance of individual i in period t , $\mathbf{x}'_i$ is a vector of individual exogenous characteristics that are constant over time, $\boldsymbol{\beta}$ are the associated parameters, and ϵ_{it} is the disturbance term reflecting unobserved heterogeneity. The expression $p_{it}^2/2$ represents the convex cost of effort. Assuming an interior solution, the first-order condition is linear and given by

$$p_{it} = \mathbf{x}'_i \boldsymbol{\beta} + \epsilon_{it}. \quad (2)$$

Second, conditional on choosing the Manual mode, we assume that the utility function is also quadratic and given by

$$V_{it} = (\mathbf{x}'_i \tilde{\boldsymbol{\beta}} + \tilde{\epsilon}_{it}) p_{it} - \frac{p_{it}^2}{2} + (\mathbf{x}'_i \boldsymbol{\alpha} + \mu_{it}) l_{it} - \frac{l_{it}^2}{2} + (\mathbf{x}'_i \boldsymbol{\kappa} + \psi_i). \quad (3)$$

The first two terms of the right-hand side of (3) reflect the sub-utility associated with the actual performance and its effort cost like in Eq. (1). However, the impact of an increase in performance might differ according to the chosen mode, because of the cost of the cognitive effort requested to compute payoffs manually in the Manual mode. This impact may also vary by individuals' characteristics ($\tilde{\boldsymbol{\beta}} \neq \boldsymbol{\beta}$ and $\tilde{\epsilon}_{it} \neq \epsilon_{it}$).

The last three terms characterize the sub-utility of choosing the Manual mode and of lying. The variable l_{it} denotes the size of lies for individual i in period t , defined as the difference between r_{it} , the reported performance, and p_{it} , the actual performance. $\boldsymbol{\alpha}$ is a vector of parameters associated with $\mathbf{x}_i$, and μ_{it} reflects unobservable heterogeneity. Over-reporting performance is the source of additional payoffs but may be morally costly because it violates a moral norm. We assume a convex cost of cheating ($= l_{it}^2/2$). The last term reflects the net fixed benefits (or costs) of choosing the Manual mode. The benefits might be from the value of controlling the computation of payoffs, from retaining the option of lying or not, or from signaling to the experimenter that an individual's type is moral if he chooses the Manual mode but does not lie. The costs might be from having to resist the temptation to lie if one's type is moral, or it might represent the reputational cost of choosing a mode that allows lying. The net benefits depend on the individual's observable and unobservable characteristics ($= \mathbf{x}'_i \boldsymbol{\kappa} + \psi_i$). Note that the functional form of the utility function given by Eq. (3) imposes separability between the actual performance p_{it} and the size of lies l_{it} .

Consistent with the literature on lying, we assume that the individuals who choose the Manual mode never under-report their performance. Therefore, we consider under-reports occurring in the experiment as mistakes. Therefore, the inequality constraint, $l_{it} \geq 0$, is assumed to hold. We suppose that $p_{it} > 0$, that is, even under the Manual mode, individuals are willing to put effort into performing the task. Indeed, the literature has shown that full lying is rare because of self-concept maintenance or perceived cheating aversion (e.g., Fischbacher and Föllmi-Heusi, 2013; Dufwenberg and Dufwenberg, 2018; Gneezy et al., 2018; Abeler et al., 2019), and our results are consistent with this assumption (full lying, that is, exerting no effort but reporting 20 grids solved, only occurred 5 times in part 1). The first-order Kuhn-Tucker conditions associated with maximizing (3) for p_{it} and l_{it} can be written as

$$p_{it} = \mathbf{x}'_i \tilde{\boldsymbol{\beta}} + \tilde{\epsilon}_{it}, \quad (4)$$

$$l_{it}^* = \mathbf{x}'_i \boldsymbol{\alpha} + \mu_{it}, \quad (5)$$

$$\text{with } l_{it} = \mathbf{I}(l_{it}^* > 0) l_{it}^*, \quad (6)$$

where l_{it}^* is the latent variable associated with the size of lies, and the indicator function $\mathbf{I}(\cdot)$ takes value 1 when $l_{it}^* > 0$ and 0 otherwise. Due to the assumption that $p_{it} > 0$, one has $0 \leq l_{it} < r_{it}$.

Finally, using Eq. (3) with $p_{it} = l_{it} = 0$, we obtain the optimal mode choice from the following equations:

$$V_{i0} = \mathbf{x}'_i \boldsymbol{\kappa} + \psi_i, \quad (7)$$

$$\text{with } M_i = \mathbf{I}(V_{i0} > 0), \quad (8)$$

where V_{i0} can be interpreted as the (latent) utility of choosing the Manual mode net of the reserve utility of choosing the Automatic mode (normalized to 0). In Eq. (8), M_i is a dummy variable equal to 1 if $V_{i0} > 0$, in which case the individual chooses the Manual mode, and equal to 0 otherwise, in which case he chooses the Automatic mode.

4.2. Generalized model with social interactions

We now generalize the model, in turn with exogenous and endogenous social interactions.

4.2.1. Exogenous matching

The individuals' utility now depends on their actual performance, their lies (if any), and the actual or reported performance of their peers. Specifically, we use a model of conformity where the individuals' social sub-utility depends on the gap between their behavior and the behavior of their peers (e.g., Boucher and Fortin, 2016; Ushchev and Zenou, 2020), which results in a linear-in-means model. Observing peers' performance under a given mode generates a reference point: to increase utility, individuals may conform to their peers to reduce the distance with them. Next, we consider each utility function, conditional on the chosen mode.

1. The individual chooses the Automatic mode

In this case, the individual's utility function is given by

$$U_{it}^{EXO} = U_{it} - (1 - m_{-i}) \frac{\lambda_1}{2} (p_{it} - \bar{p}_{-it})^2 - m_{-i} \frac{\lambda_2}{2} (p_{it} - \bar{p}_{-it}^e)^2, \quad (9)$$

where the first expression on the right-hand side is the individual's *private* sub-utility function, the sum of the two other expressions represents his *social* sub-utility function, and m_{-i} is an exogenous binary variable that is equal to 0 if peers chose the Automatic mode and 1 if they chose the Manual mode.

In the presence of (weak) conformity, $\lambda_1 \geq 0$: the individual's social sub-utility decreases when peers chose Automatic, and the (Euclidian) distance between the individual's actual performance and the average actual performance of his peers becomes larger. Moreover, $\lambda_2 \geq 0$: the individual's social sub-utility decreases when peers chose Manual, and the distance between his actual performance and the average *expected* performance of his peers, $\bar{p}_{-it}^e$, increases.

The latter variable is not perfectly known by the individual who only receives an imprecise signal of the average actual performance when peers chose Manual, that is, their average reported performance. We assume that an individual forms his expectation of $\bar{p}_{-it}$ as an increasing function of the average performance reported by his peers: $\bar{p}_{-it}^e = \phi(\bar{r}_{-it})$, with $\phi(0) = 0$, $\phi'(\bar{r}_{-it}) \geq 0$, and $0 \leq \bar{p}_{-it}^e \leq \bar{r}_{-it}$. As a first-order approximation, we assume that $\bar{p}_{-it}^e = \delta \bar{r}_{-it}$, where δ is the expectation coefficient, with $0 \leq \delta \leq 1$. Assuming that individuals form accurate beliefs on the actual performance of their peers, the difference between λ_1 and λ_2 indicates whether individuals put a different weight on their peers depending on the origin of their performance when they choose their level of effort.

The first-order condition yields the following best response function:

$$p_{it} = \frac{\mathbf{x}'_i \boldsymbol{\beta} + (1 - m_{-i}) \lambda_1 \bar{p}_{-it} + m_{-i} \lambda_2 \delta \bar{r}_{-it} + \epsilon_{it}}{\Delta_1}, \quad (10)$$

where $\Delta_1 = 1 + (1 - m_{-i}) \lambda_1 + m_{-i} \lambda_2 \geq 1$. Eq. (10) indicates that in the presence of conformity, the average peers' (actual or reported) performance has a positive impact on the individual's actual performance. The expression $\lambda_1/(1 + \lambda_1)$ represents the effect of the average actual performance from peers who chose the Automatic mode ($m_{-i} = 0$); $\lambda_2 \delta/(1 + \lambda_2)$ represents the effect of the reported performance from peers who chose the Manual mode ($m_{-i} = 1$).

2. The individual chooses the Manual mode

In this case, the individual's payoff depends on his reported performance, r_{it} , composed of his actual performance, p_{it} , and his lies of size l_{it} . We assume that both components can be influenced by peers. The individual's utility function is given by

$$V_{it}^{EXO} = V_{it} - (1 - m_{-i}) \frac{\lambda_3}{2} (p_{it} - \bar{p}_{-it})^2 - (1 - m_{-i}) \frac{\lambda_4}{2} l_{it}^2 - m_{-i} \frac{\lambda_5}{2} (p_{it} - \bar{p}_{-it}^e)^2 - m_{-i} \frac{\lambda_6}{2} (l_{it} - \bar{l}_{-it}^e)^2, \quad (11)$$

where the first expression on the right-hand side is the individual's *private* sub-utility function, and the sum of the four other expressions denotes his *social* sub-utility function.

If the individual is matched with peers who chose the Automatic mode, he observes their average actual performance, $\bar{p}_{-it}$. In the presence of conformity, he can increase his social sub-utility by reducing the gap between his actual performance and the average actual performance of his peers by adjusting his effort level (an effect denoted $\lambda_3 \geq 0$) or by reducing the size of his lies (an effect denoted $\lambda_4 \geq 0$). If the individual is matched with peers who chose the Manual mode, he can increase his social sub-utility by reducing the gap between his actual performance and the mean expected actual performance of his peers. This effect is denoted $\lambda_5 \geq 0$. He can also increase his social sub-utility by reducing the gap between his lies

and the average expected lies of his peers. This effect is denoted $\lambda_6 \geq 0$. In other words, in the presence of conformity, the individual's lying behavior is not influenced by the expected actual performance of peers but by their expected misreporting behavior.

The first-order Kuhn-Tucker conditions of (11) yield the following best response functions:

$$p_{it} = \frac{\mathbf{x}_i' \tilde{\beta} + (1 - m_{-i})\lambda_3 \bar{p}_{-it} + m_{-i}\lambda_5 \delta \bar{r}_{-it} + \tilde{\epsilon}_{it}}{\Delta_2}, \quad (12)$$

$$l_{it}^* = \frac{\mathbf{x}_i' \alpha + m_{-i}\lambda_6(1 - \delta)\bar{r}_{-it} + \mu_{it}}{\Delta_3}, \quad (13)$$

$$= \frac{\mathbf{x}_i' \alpha + \mu_{it}}{1 + \lambda_4}, \quad \text{when } m_{-i} = 0 \quad (13')$$

$$= \frac{\mathbf{x}_i' \alpha + \lambda_6(1 - \delta)\bar{r}_{-it} + \mu_{it}}{1 + \lambda_6}, \quad \text{when } m_{-i} = 1 \quad (13'')$$

$$\text{with } l_{it} = \mathbf{I}(l_{it}^* > 0) l_{it}^*, \quad (14)$$

where $\Delta_2 = 1 + (1 - m_{-i})\lambda_3 + m_{-i}\lambda_5 \geq 1$ and $\Delta_3 = 1 + (1 - m_{-i})\lambda_4 + m_{-i}\lambda_6 \geq 1$. Eq. (12) indicates that an increase in the average peers' performance (actual or reported) observed by individual i will increase his actual performance. The expressions $\lambda_3/(1 + \lambda_3)$ and $\lambda_5\delta/(1 + \lambda_5)$ represent the effects of the average actual and reported performance of peers on the individual's actual performance, in the case of peers choosing the Automatic or the Manual mode, respectively.

Our model suggests that the size of lies, our primary focus, depends on the individual's characteristics and his peers' expected behavior. When the individual is matched with peers who chose the Automatic mode ($m_{-i} = 0$), he knows that they did not lie. In such a case, our conformity model predicts that the size of lies does not depend on peers' performance, but only on being matched with peers who behaved honestly (due to the impact of $1 + \lambda_4$ on l_{it}^* in Eq. (13')).¹¹

When the individual is matched with peers who chose the Manual mode ($m_{-i} = 1$), conformity implies that he will increase the size of his lies when their reported performance becomes higher. We assume that reason for this is that the individual forms an expectation of his peers' lying behavior as a proportion of their average reported performance. In Eq. (13''), the expression $\lambda_6(1 - \delta)/(1 + \lambda_6) \geq 0$ represents the effect of peers' average reported performance under the Manual mode on the individual's lying behavior. Note that $\lambda_6/(1 + \lambda_6)$ (or λ_6) cannot be identified because the expectation coefficient, δ , is not identifiable. This can be understood intuitively: if someone observes that other people who selected an environment that allows them to lie report a higher performance, he can infer that these people probably have lied. This may change the perception of the norm and work as a self-excusing justification for individuals over-reporting their performance.

Proposition 1. (Peer Effects on Lying) (a) The lying behavior of an individual who chose the Manual mode is not influenced by his peers performance when they have chosen the Automatic mode ($m_{-i} = 0$).

(b) By contrast, when peers have chosen the Manual mode ($m_{-i} = 1$), and in the presence of preferences for conformity, lies increase with peers' reported performance.

Proof. (a) Proposition 1a follows from our assumption that the preferences of an individual who chose the Manual mode and interacts with peers is given by Eq. (11), which supposes, among others, separability between $(p_{it} - \bar{p}_{-it})$ and l_{it} . From this assumption, the FOC associated with the size of lies yields Eqs. (13') and (14) when $m_{-i} = 0$, that is, when peers chose the Automatic mode. This implies Proposition 1a, as $\bar{p}_{-it}$ does not appear on the right-hand side of Eq. (13'). (b) Since $\lambda_6(1 - \delta)/(1 + \lambda_6) \geq 0$ in Eq. (13''), the effect of the mean performance reported by peers who selected the Manual mode ($m_{-i} = 1$) is positive on the individual's size of lies. $\square$

4.2.2. Endogenous matching

We now consider endogenous social interactions. At the individual level, link formation is assumed to be homophilous when an individual selects peers who chose the same mode as him (i.e., signaling similar values). However, at the researcher level, a component of the individual's preferences is random because of unobservable heterogeneity. We assume homophilous link formation when the individual's likelihood of choosing peers who selected the same mode is higher than chance. Two main reasons can influence such assortative matching.

The first reason is the activation of a feeling of similarity (for preference-based models of group identity, see Akerlof and Kranton (2000); for selection based on weak signals, see Efferson et al. (2008)). Formally, suppose an individual i does not like to compute his payoffs manually and values honesty. Therefore, he chooses the Automatic mode. This variable is a predetermined variable in the model. Now, he must select a pair of peers based on the information on their chosen mode

¹¹ Note that if peer effects were driven by inequality aversion, individuals could adjust their lying behavior to the actual performance of their peers who chose the Automatic mode. This is different from our model of conformity.

and their average performance in part 1. Conditional on his choice of the Automatic mode, his (latent) utility of selecting peers who also chose the Automatic mode is assumed to be given by

$$a_i^* = \mathbf{a}_1 \mathbf{x}_i' + a_2 P_i^A + a_3 P_i^M + a_4 |P_i^A - P_i^M| + \xi_i, \quad (15)$$

where P_i^A and P_i^M are the average actual performance in part 1 of the peers who chose the Automatic mode and the average reported performance of the peers who chose the Manual mode, respectively. ξ_i is the random term. The expression in absolute value introduces a potential nonlinearity in the effects of the peers' actual and reported performance. Assuming that his reserve utility of selecting the peers who chose the Manual mode is 0, the decision to select the peers who chose Automatic is given by

$$a_i = \mathbf{I}(a_i^* > 0), \quad (16)$$

where a_i is a dummy variable equal to 1 if the individual who chose the Automatic mode selects the peers who made the same choice and equal to 0 otherwise. Next, we suppose that Eqs. (15) and (16) are estimated by using a linear probability model. We also assume, for notational simplicity, that the mean of all the right-hand side variables is normalized to 0 over the relevant population. Then, a majority of the population who chose the Automatic mode is homophilous if, at the average of the explanatory variables (normalized to 0), the constant of the regression is larger than 0.5: $E[E(a_i)|\mathbf{x}_i, P_i^A, P_i^M] > 0.5$. Not rejecting this hypothesis indicates that a majority of participants who chose Automatic preferred to select peers with the same mode choice.

Although a taste for similarity might also apply to individuals who chose the Manual mode, a second reason for homophilous link formation concerns the individuals who selected the Manual mode and lied in part 1. These individuals may be less attracted by peers who selected a different mode because being exposed to peers who are less likely to have lied may recall the injunctive moral norm of honesty and make one's misconduct more salient. This captures the idea of willful ignorance of inconvenient information. By contrast, if peers' behavior is perceived as a signal of an empirical norm, observing others reporting a higher performance might help liars convince themselves that over-reporting is what the majority of others probably do. It may offer to liars an excuse to justify unethical behavior, whereas individuals who chose the Manual mode but do not lie do not need such arguments. We thus expect more homophilous link formation among liars.

Conditional on choosing Manual, given the information on the mode and performance of peers, the (sub-)utility of player i who selects peers who chose Manual is

$$m_i^* = \mathbf{x}_i' \mathbf{b}_1 + b_2 P_i^A + b_3 P_i^M + b_4 |P_i^A - P_i^M| + b_5 D_i^D + \eta_i, \quad (17)$$

where D_i^D is a predetermined dummy variable equal to one if the individual lied in part 1 and zero otherwise. Consistent with our aforementioned argument, we assume that $b_5 \geq 0$, that is, if the individual chose the Manual mode and lied, his utility increases when he selects peers who also chose the Manual mode. Assuming that his reserve utility when selecting peers who chose the Automatic mode is zero, the decision to select the peers who chose the Manual mode is given by

$$m_i = \mathbf{I}(m_i^* > 0). \quad (18)$$

Again, we assume that Eqs. (17) and (18) are estimated by using a linear probability model and that the regressors are normalized so that their average over the relevant population is 0. Then, a majority of the Manual population in ENDO will form homophilous links if, at the average of the explanatory variables ($= \mathbf{0}$), the constant of the regression is larger than 0.5: $E[E(m_i)|\mathbf{x}_i, P_i^A, P_i^M, D_i^D] > 0.5$. This leads to Proposition 2.

Proposition 2. (Homophilous Link Formation) (a) In ENDO, individuals prefer to be matched with participants who have selected the same mode as the individuals selected. (b) *Ceteris paribus*, assortative matching is more likely among individuals who chose the Manual mode and lied than among the other individuals who chose the Manual mode.

Proof. Proof: (a) This is a consequence of the hypothesis that $E[E(a_i)|\mathbf{x}_i, P_i^A, P_i^M] > 0.5$ for the participants who chose the Automatic mode and $E[E(m_i)|\mathbf{x}_i, P_i^A, P_i^M, D_i^D] > 0.5$ for those who chose the Manual mode. $\square$

Part (b) of the proposal follows from the assumption that $b_5 \geq 0$ in Eq. (17).

4.2.3. Peer effects, self-Selection, and link strength effect

Homophilous link formation may induce an overestimation of the causal peer effects in the ENDO treatment because of a self-selection bias. This might occur if individuals report a similar performance as their peers not because of causal peer effects on behavior, but because they have selected peers with a similar morality. In addition, the possibility given to the participants to select their peers may increase the size of peer effects on lying. This may occur if individuals value more the examples of peers precisely because they have actively selected them. This effect is not the source of a bias, but its presence indicates that the causal peer effects might depend on the strength of the links between the participants and their peers.

If present, these two effects would move the estimation of peer effects in the same direction. Thus, a simple strategy is pooling the data of individuals who observed the performance of peers who chose the Manual mode and testing whether the significance and magnitude of the peers' reported performance on the size of lies differ in EXO and ENDO. This offers a joint test of the two effects. If the influence of peers' reported performance on the size of lies is statistically indistinguishable in the two treatments, we can reject both the self-selection bias and the link strength effect.

In our econometric strategy, we assume that, under the null hypothesis, these two effects do not influence the causal peer effects. Therefore, m_i (which identifies the mode of selected peers) is exogenous, that is, uncorrelated with the disturbance term associated with the lying behavior, even in ENDO. Thus, in the absence of a link strength effect, the lying behavior equations for an individual whose peers have chosen Manual (Eqs. (13'') and (14)) should be the same in both EXO and ENDO. To jointly test the null hypothesis of no self-selection bias and no link strength effect, we can write the following general model that generalizes Eq. (13'') for both treatments, conditional on $m_{-i} = 1$. We obtain

$$l_{it}^* = \mathbf{x}_i' \tilde{\theta} + \theta_1 (D_{EXO}) \bar{r}_{-it} + \theta_2 (D_{ENDO}) \bar{r}_{-it} + \theta_3 D_{ENDO} + \zeta_{it}, \quad (19)$$

where D_{ENDO} is a dummy equal to 0 for the EXO treatment and 1 for the ENDO treatment, $D_{EXO} = 1 - D_{ENDO}$, and where the θ s are parameters to be estimated.

Under the null hypothesis, $E(\zeta|\cdot) = 0$ and $\theta_1 = \theta_2$.

θ_3 controls for the effect of ENDO on the intercept.

Eq. (19) can be rewritten as

$$l_{it}^* = \mathbf{x}_i' \tilde{\theta} + \theta_1 \bar{r}_{-it} + \theta_4 (D_{ENDO}) \bar{r}_{-it} + \theta_3 D_{ENDO} + \zeta_{it}, \quad (20)$$

where $\theta_4 \geq 0$ captures the presence of a self-selection bias in the peer effects or a link strength effect. Then, a one-tailed test of whether $\theta_4 (= \theta_2 - \theta_1)$ is 0 against the alternative hypothesis that it is larger than 0 is sufficient to determine whether the endogenous choice of peers affects the causal peer effects. A one-tailed test is sufficient because there is no reason to expect that $\theta_2 < \theta_1$. This analysis leads to Proposition 3:

Proposition 3. (Peer Effects, Self-Selection Bias, and Link Strength) In ENDO, we assume the presence of homophilous link formation (Proposition 2). The estimated effect of the average reported performance of peers who chose the Manual mode on individuals' lies may be biased upward because of self-selection or increased by a link strength effect generated by the selection of peers. A comparison between the estimators under EXO and ENDO provides a joint test of the absence of both these features.

5. Results

In this section we highlight four points.¹² First, we study the selection of the mode of evaluation and lying behavior in isolation. Second, we test Proposition 1 by analyzing the effect of social information on lying behavior in EXO. Third, we test Proposition 2 by exploring the determinants of the choice of peers in ENDO. Finally, we test Proposition 3 by examining whether the selection of peers influences the measurement of peer effects through a self-selection bias or a link strength effect.

5.1. Mode choice and lying behavior in isolation

We start by stating our first result.

Result 1. (Mode Choice and Lying Behavior in Isolation) Almost half of the participants chose the Manual mode that allowed them to lie and a slight majority of them actually over-reported their performance when they worked in isolation.

Support for Result 1. 47.97% of the participants chose the Manual mode (43.06% in the Baseline, 51.05% in EXO, and 47.40% in ENDO). No pairwise comparison is significant (t -tests, $p = .270$ for EXO vs. Baseline, $p = .543$ for ENDO vs. Baseline, and $p = .532$ for ENDO vs. EXO), which was expected because part 1 was identical across treatments. Based on Eqs. (7) and (8), Table A.3 in the Appendix displays the marginal effects of Probit regressions on the choice of the Manual mode. In model (1), the independent variables include individual characteristics; in model (2), EXO and ENDO dummies are added as regressors, as a sanity test. The choice of the mode does not depend on observed individual characteristics in any model (none are significant at conventional levels), and there are no significant treatment effects.

Table 1 summarizes the means and standard deviations of the actual performance, the reported performance, and the size of lies per individual and per period in part 1, according to the chosen mode and by treatment. It also displays the percentage of participants who over-reported their performance at least once in part 1 and the percentage of "Manual-Dishonest" participants. Manual-Dishonest participants are defined as individuals whose average difference between the reported and actual performances in part 1 was higher than 1 per period, which allows for small calculation mistakes; otherwise, participants are classified as "Manual-Honest". In our analysis, we define a lie as any positive difference between the actual and reported performance. Some misreporting could occur because of mistakes but the estimated frequency of mistakes is very low.¹³

¹² We do not extensively discuss the actual performance in the text because it is not the main aim of this study. The effects of explanatory variables on this outcome are very minor, as we explain next.

¹³ We estimate the frequency of mistakes by looking at the cases in which the actual performance was under-reported, – which, we assumed, is not lying –, as mistakes should be symmetrical. This happened in only 35 cases out of 1770 (1.98% of the reports). In 19 under-reports the difference with the actual performance was equal to 1; in 12 under-reports it was equal to 2 and differences higher than 2 occurred only four times; the average size of

Table 1
Performance and lies in Part 1, by mode and treatment.

Treatment	Baseline		EXO		ENDO		ALL	
Chosen mode	Auto	Manual	Auto	Manual	Auto	Manual	Auto	Manual
Mean actual performance	12.87 (3.60)	13.79 (4.85)	14.46 (3.56)	12.76 (4.93)	13.25 (3.17)	13.21 (3.96)	13.61 (3.45)	13.12 (4.53)
Mean reported performance	-	21.05 (9.11)	-	20.83 (8.38)	-	18.73 (6.37)	-	20.00 (7.78)
Mean lie	-	7.26 (11.33)	-	8.08 (10.80)	-	5.51 (7.72)	-	6.88 (9.76)
% Over-reporters	-	61.29%	-	65.75%	-	67.12%	-	65.54%
% Manual-Dishonest	-	51.61%	-	53.42%	-	52.05%	-	52.54%
N	41	31	70	73	81	73	192	177
Percentage	56.94	43.06	48.95	51.05	52.60	47.40	52.03	47.97

Notes: The table reports mean values. Standard deviations are reported in parentheses below mean values. "% Over-reporters" corresponds to the extensive margin (percentage of participants who over-reported at least once during part 1). "% Manual-Dish." corresponds to the percentage of participants whose reports deviated from their actual performance by more than 1 on average per period. The mean lie in a period is the mean difference between the reported and the actual performance when participants chose the Manual mode. Performance is expressed as the number of points earned in a period.

Table 1 shows that not all participants who chose the Manual mode lied. Overall, 65.54% over-reported their performance at least once during part 1 (61.29% in the Baseline, 65.75% in EXO, and 67.12% in ENDO). No pairwise comparison is significant in *t*-tests ($p = .667$ for EXO vs. Baseline, $p = .572$ for ENDO vs. Baseline, and $p = .862$ for ENDO vs. EXO). We classify 52.54% of the participants who chose the Manual mode as Manual-Dishonest participants overall (51.61% in the Baseline, 53.42% in EXO, and 52.05% in ENDO).

No pairwise comparison is significant ($p = .867$ for EXO vs. Baseline, $p = .967$ for ENDO vs. Baseline, and $p = .869$ for ENDO vs. EXO). Most participants decided to be honest (either by choosing the Automatic mode or choosing the Manual mode and reporting honestly). This suggests the presence of a strong social norm of honesty. It is consistent with the high prevalence of honesty found in the previous literature on lying (see, e.g., the meta-analysis of Abeler et al., 2019). In addition, it is possible that our two-step process (first, choosing a mode and then, performing a task where cheating is possible) increased participants' reflection and influenced lying downward. That at least 34.46% of those who chose Manual did not lie reveals that non-pecuniary benefits compensated the cognitive cost of having to calculate one's performance.¹⁴

Our model assumes that the actual performance is strictly positive ($p_{it} > 0$). Data confirm that individuals in the Manual mode lied only partially. A very small minority lied to the full extent: in part 1, five individuals (2.82%) reported 40 points in each period (corresponding to 20 grids paid at the rate of 2 points since individuals could lie both on the piece rate and the number of correct answers), and five decisions (from different individuals) correspond to the maximum possible lie (reporting 40 when the performance was 0). Partial lying suggests the presence of self-image concerns. When participants over-reported, on average they inflated performance by 52.44% overall (average actual performance: 13.12; average lie: 6.88), by 52.65% in the Baseline, by 63.33% in EXO, and by 41.71% in ENDO. If we restrict our scope to the Manual-Dishonest individuals, they inflated their performance by 113% overall (average actual performance: 11.50; average lie: 13.02), by 115% in the Baseline, by 135% in EXO, and by 91% in ENDO. More generally, none of the differences across treatments in Table 1 are significant according to Mann-Whitney tests (MW, hereafter), except that the actual performance of participants who chose the Automatic mode is higher in EXO than in the other treatments (MW, EXO vs. Baseline, $p = .041$, and EXO vs. ENDO, $p = .021$).¹⁵

When pooling the treatments, we find that the average actual performance is not significantly different for the participants who chose the Manual mode and for those who chose the Automatic mode (MW, $p = 0.798$). This is also true when considering each treatment separately, except for EXO, in which participants who chose Automatic performed marginally significantly better than those who chose Manual (MW, $p = 0.072$). Considering only the Manual-Dishonest participants reveals a significant difference in actual performance with the participants who chose the Automatic mode (MW, $p = 0.002$), suggesting that for these individuals lying partially substituted for effort. However, their actual performance was rarely 0

under-reports was 1.71. These under-reports were present in any part, without significant differences across treatments. Since negative mistakes were rare and small in magnitude, we conclude that the vast majority of over-reports were intentional lies. Nevertheless, to account for these mistakes, we defined a participant as Dishonest if the size of his lies per period in part 1 was greater than 1 on average.

¹⁴ One reason given by these participants in the final questionnaire was that they wanted to have more control over their earnings (perhaps being anxious about their future performance, although their mean performance in part 1 was significantly higher than those who chose Automatic, $p = .003$) and they preferred to calculate themselves. The need for control has been identified as one of the five core social motives that propel human behavior Fiske (2018). The need for control is also closely related to another social motive: trusting (here, trusting the calculation of the performance by the program). Some of these participants were also willing to test their honesty. Another possible reason –although not mentioned in the questionnaire– is signalling their honesty to the experimenter by exposing themselves to a context where lying was possible and not lying.

¹⁵ The reported non-parametric statistics are based on averaged measures per individual such that each individual in each part gives one independent observation. Except if specified otherwise, the reported tests are two-tailed.

Table 2
Peer effects on the size of lies under the manual mode, EXO, Part 2.

Dep. var.: Size of lies	(1) Manual All peers	(2) Manual Auto peers	(3) Manual Manual peers	(4) Manual-H Manual peers	(5) Manual-D Manual peers
(AUTO) Peers' actual perf.	–0.005 (0.123)	–0.062 (0.110)	– –	– –	– –
(MANUAL) Peers' reported perf.	0.194** (0.076)	– –	0.229* (0.117)	–0.146 (0.151)	0.432*** (0.153)
Individual characteristics	Yes	Yes	Yes	Yes	Yes
<i>N</i>	365	175	190	65	125
Left-censored obs.	144	87	57	50	7
Log-likelihood	–761.18	–297.83	–458.46	–62.60	–376.97
Log-likelihood test $\lambda_4 = \lambda_6$ (p-value)	0.002***				

Notes: This table reports the average marginal effects over time and individuals from RE Tobit regressions with standard errors in parentheses. The dependent variable is the size of lies in each period of part 2 in EXO. There are 5 observations per individual. (AUTO) Peers' actual perf. is the mean actual performance of the two peers who chose Automatic. (MANUAL) Peers' reported perf. is the mean reported performance of the two peers who chose Manual (its coefficient is the marginal effect corresponding to $\lambda_6(1 - \delta)/(1 + \lambda_6)$ in our model). The individual characteristics include age in years, gender, educational achievement (from 0 to 9), student status, studying business, monthly expenses, and a binary variable for a first participation in an experiment. Models with Manual-H or Manual-D, resp., include only the participants who chose Manual and were classified as honest or dishonest in part 1. Models with Auto peers (Manual, resp.) include only the participants who were matched with peers who chose Automatic (Manual). * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$.

(4.3% of the observations; see Figure A.5 in Appendix). We also find some heterogeneity in lying by effort, as the average size of lies is higher in the first quintile of the average actual performance distribution but very similar in the other quintiles (see Figure A.6 in Appendix).

Finally, we estimate Eqs. (5) and (6) to analyze the impact of sociodemographic variables on the size of lies in part 1. Table A.4 in the Appendix displays the marginal effects of random-effect (RE) Tobit regressions in which the dependent variable is the (latent) size of lies in each of the five periods of part 1.

Although the estimators may be subject to a selection bias, the regressions on the mode choice reported in Table A.3 in Appendix suggest that it is not likely to be the case, because no observable demographic variables are significant at the 5% level.

5.2. Peer effects on lying behavior

We now consider part 2. Descriptive statistics on actual performance, reported performance, and size of lies are summarized in Table A.5 in Appendix. According to Proposition 1, the size of an individual's lies is not affected by the average actual performance of peers who chose the Automatic mode and thus could not lie, but it can be influenced by the average reported performance of peers who chose the Manual mode (taken as an approximation of their expected lies; Eqs. (13) and (14)). To test this proposition, we study the effects of peers' performance on the size of lies, controlling for individual characteristics, in EXO, where a self-selection bias can be ruled out by design. Our second result supports Proposition 1 and is summarized as follows:

Result 2. (Peer Effects on Lying) In EXO, the size of lies of individuals who chose the Manual mode was not affected by the performance of peers who selected the Automatic mode, whereas it increased with a higher reported performance of peers that selected the Manual mode. Such conformity with this category of peers is found among individuals who were lying in part 1 (Manual-Dishonest individuals) but not among those who did not lie in part 1 (Manual-Honest individuals).

Support for Result 2. Table 2 displays the marginal effects of RE Tobit regressions in which the dependent variable is the size of lies in each period of part 2, conditional on choosing the Manual mode. Model (1) in Table 2 pools the participants from EXO who selected the Manual mode and provides estimates of Eq. (13) augmented for Auto peers' actual performance. From Eqs. (13') and (13''), one easily shows that this model restricts that the marginal effects of individual characteristics are not affected by the mode chosen by peers (that is, $\lambda_4 = \lambda_6$, which is tested next). Model (2) (respectively, model (3)) is estimated based on participants who selected the Manual mode and were matched with peers who chose the Automatic mode (respectively, peers who chose the Manual mode). Thus, models (2) and (3) allow us to test Eqs. (13') and (13'').

Models (4) and (5) consider participants who selected the Manual mode and were matched with peers who made the same choice, but they are estimated based on the sub-samples of Manual-Honest participants ("Manual-H" in the Table) and Manual-Dishonest participants ("Manual-D"), respectively, as defined by their behavior in part 1. The most important independent variable is the actual or reported performance of peers, according to their chosen mode. These regressions control for the participants' characteristics; none of them are significant at standard levels ($\leq 5\%$ level).

Consistent with the first part of [Proposition 1](#), in models (1) and (2) the size of lies is not influenced by the peers' actual performance under the Automatic mode.¹⁶ By contrast, and consistent with the second part of [Proposition 1](#), the higher the observed reported performance from peers who chose the Manual mode, the higher the participants' size of lies (see models (1), (3), and (5)). This significant effect is clear evidence of positive peer effects on dishonesty.¹⁷ The magnitude of these peer effects is not small. An increase of one earnings point by peers increased an individual's lie by 0.194 points in EXO (model (1)). The impact is stronger if we focus on the individuals matched with peers who chose the Manual mode (model (3)), and even stronger when the individual was dishonest in part 1 (the increase amounts to 0.432 points in model (5)).

Notably, the effect of peers vanishes when we restrict our scope to the Manual-Honest participants: in model (4), peers' performance is not significant.¹⁸ In EXO, only five participants that selected the Manual mode (7% of the Manual participants) switched from the Manual-Honest to the Manual-Dishonest category between parts, that is, their average size of lies per period became greater than 1 in part 2. Moreover, the average size of lies in part 2 of these switchers was low: they inflated their performance by 6.04 points per period on average (Manual-Dishonest participants did it by 19.63 points on average). Most of the Manual-Honest participants remained fundamentally honest despite the information received.^{19,20} The honesty norm of these participants was a powerful driver of their behavior. The immutability of their behavior in the second part of the experiment shows their strong willingness to follow the honesty norm.

Finally, model (1) imposes that the marginal effects of individual characteristics are not affected by the peers' mode (i.e., $\lambda_4 = \lambda_6$). We test this by comparing models (2) and (3) with model (1), using a log-likelihood test (see [Table 2](#)). This equality is clearly rejected (p -value=0.002), which justifies estimating models (2) and (3) separately.

5.3. Choice of peers

We now examine how participants chose their peers. According to [Proposition 2](#), in ENDO, individuals prefer to create links with peers who chose the same mode as the individuals did (a), and such a homophilous link formation should be more likely among the individuals who misreported their performance when working in isolation (b). Our results support (b) and partially (a):

Result 3. (Link formation) Only the individuals who chose the Manual mode and lied in isolation were more likely to form links with peers who made a similar choice of mode.

Support for Result 3. In ENDO, 57.79% of the participants selected peers who chose the Manual mode. This percentage differs across participants depending on their chosen mode and behavior in part 1 ([Fig. 2](#)). 55.55% of the participants who chose Automatic selected peers who made the same choice. This percentage is not significantly different from 50% (binomial test, $p=0.374$), which rejects the presence of homophily based on a taste for similarity (in terms of mode choice). By contrast, 72.60% of the participants who chose the Manual mode formed links with peers that also chose Manual. If among this group we consider only the Manual-Dishonest participants, this percentage is even larger (81.58%), while it is 62.86% for the Manual-Honest participants. The percentage is significantly higher than a random choice of peers for Manual-Dishonest participants (binomial test, $p < .001$), but not for Manual-Honest participants ($p = .175$). Liars in isolation were willing to connect to individuals who were likely to behave as they did.

[Table 3](#) provides parametric tests based on [Eqs. \(15\)–\(18\)](#), using a linear probability approach. OLS regressions are performed for all participants (models (1) and (2)), and for those who chose Automatic (model (3)), or Manual (model (4)), separately; the latter are decomposed into Manual-Honest (model (5)) and Manual-Dishonest participants (model (6)). The dependent variable is equal to 1 when the individual selected two peers that chose the same mode as him, and 0 otherwise. The independent variables include the information provided to the participants before they made their choice of peers, namely, the average actual performance (P_i^A in the theoretical model) and reported performance (P_i^M in the model) in part 1 of the two pairs of peers that were presented to the participant. We include the absolute difference between these

¹⁶ This is also true when estimating model (2) separately for the Manual-H participants and the Manual-D participants who were matched with peers who chose the Automatic mode. The coefficients are 0.004 with a standard error of 0.047 for Manual-H and -0.199 with a standard error of 0.224 for Manual-D. We omit these two regressions in [Table 2](#) for the sake of concision.

¹⁷ Note that individuals did not lie to equalize their reported performance to that of their peers. Indeed, lying depends on a variety of determinants and not only on peers' performance. Equalizing performance represented only 4.18% of the observations. Participants reported less than their peers' average performance in 52.09% of the observations and more in 43.74% of the observations. In EXO, they reported less than their peers in 46.32% (more in 49.47%) of the observations; in ENDO, the respective percentages are 56.23% and 39.62%. These percentages are not significantly different between EXO and ENDO (MW, $p = .375$ and $p = .367$, respectively).

¹⁸ As a robustness test, we replicated [Table 2](#) with a stricter definition of Manual-Dishonest participants. In [Table A.7](#) in the Appendix, an individual was classified as Manual-Dishonest if he reported on average more than 2 more points than his actual performance per period in part 1 (instead of more than 1 in [Table 2](#)), and as Manual-Honest otherwise. The results remained qualitatively similar.

¹⁹ In ENDO, we found qualitatively the same results: only eight participants switched from the Manual-Honest to the Manual-Dishonest category, and the switchers inflated their performance by 3.5 points on average per period, while Manual-D participants inflated their performance by 17.94 points.

²⁰ We also analyzed the influence of peers' performance on the participants' actual performance, using the same models like those of [Table 2](#) but in which the dependent variable was the individual's actual performance. [Table A.9](#) (Table A.10, resp.) in Appendix reveals that peers' actual or reported performance had no significant influence on the performance of participants who chose Automatic (Manual, resp.) in any treatment. Peer effects did not influence actual effort.




Fig. 2. Choice of peers according to the chosen mode and lying behavior in Part 1.

Notes. The figure displays the share of peers with their mode of evaluation for three categories of participants in ENDO: participants who selected the Automatic mode (left bar), participants who selected the Manual mode but did not lie in part 1 (Manual-Honest, middle bar), and participants who selected the Manual mode and lied in part 1 (Manual-Dishonest, right bar).

Table 3
Homophilous link formation in the ENDO treatment.

Dep. var.: Choice of peers with the same mode	ALL (1)	ALL (2)	AUTO (3)	MANUAL (4)	MANUAL-H (5)	MANUAL-D (6)
Constant	0.636*** (0.039)	0.636*** (0.039)	0.556*** (0.055)	0.726*** (0.053)	0.629*** (0.088)	0.816*** (0.066)
Average performance of potential AUTO peers	0.030 (0.025)	0.033 (0.025)	0.074** (0.035)	-0.027 (0.037)	-0.014 (0.064)	-0.060 (0.065)
Average performance of potential MANUAL peers	-0.019 (0.026)	-0.019 (0.026)	-0.050 (0.036)	0.022 (0.039)	0.011 (0.063)	0.037 (0.065)
Diff (absolute) between performance of potential AUTO vs. MANUAL peers	0.030 (0.028)	0.028 (0.027)	0.050 (0.038)	-0.009 (0.041)	0.001 (0.069)	-0.034 (0.068)
Manual-D	-	0.231** (0.099)	-	-	-	-
Manual-H	-	0.041 (0.101)	-	-	-	-
Individual characteristics	Yes	Yes	Yes	Yes	Yes	Yes
One-tailed Homo.Test:						
Const.=0.50 (p-value)	0.0003***	0.0003***	0.159	0.00005***	0.079*	0.00005***
N	154	154	81	73	35	38
R ²	0.049	0.085	0.133	0.112	0.202	0.218

Notes: The table reports the coefficients from OLS regressions with standard errors in parentheses. The dependent variable is a dummy variable equal to 1 if the participant selected a pair of peers who chose the same mode as himself, and 0 otherwise. There is one observation per individual. All independent variables are net of their mean in the regressions. *Average performance of potential AUTO peers* (respectively, *Average performance of potential MANUAL peers*) are the average actual (respectively, reported) performance in part 1 of the peers who chose the Automatic (respectively, Manual) mode. *Diff (absolute) between the performance of potential AUTO vs. MANUAL peers* is the absolute difference between the last two variables described. The individual characteristics include the participant's age in years, gender, educational achievement (from 0 to 9), student status, being a student in business, monthly expenses, and a binary variable for a first participation in an economic experiment. Models with *Manual-H* or *Manual-D*, respectively, include only participants who chose the Manual mode and were classified as honest in part 1 or, respectively, as dishonest. AUTO for Automatic mode. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

two variables. We control for the same individual characteristics as in Table 2; only the variable *Student* is significant at the 5% level and positive in models (4) and (5). The observable characteristics are net of their mean; thus, the constant can be interpreted as an estimator of the average proportion of participants who prefer to be matched with similar peers. The values of the constant closely correspond to the numbers reported in Fig. 2.

When there is no homophily (an individual is indifferent between the two peers who chose the Automatic mode and the two peers who chose the Manual mode), an individual's probability to select the peers who chose the same mode as him should be close to 50%. However, Table 3 shows that, when pooling all the data (model (1)), the probability of

an individual to select the peers who chose the same mode as him reaches 63.6% (i.e., this probability increases by $63.6 - 50 = 13.6$ percentage points compared to when there is no homophily). This probability even reaches 81.6% when the individual belongs to the Manual-Dishonest category (model (5)) (i.e., an increase of $83.6 - 50 = 33.6$ percentage points compared to when there is no homophily). As mentioned above, we perform one-tailed tests where the null hypothesis is $H_0 : \text{Constant} = 0.5$ (no homophily), and the alternative hypothesis is $H_A : \text{Constant} > 0.5$ (homophily). Our tests confirm the presence of homophilous link formation for the whole population and for the participants who chose the Manual mode (models (1) and (4)). It is rejected for those who chose the Automatic mode (model (3)): this could be either because these individuals were indifferent between the two types of peers, or because the taste for similarity was dominated by the wilful avoidance of repeat actual performance comparisons. We reject the second interpretation: model (3) shows that individuals who chose the Automatic mode were more likely to select similar peers, the higher was these peers' performance; moreover, the mean performance in part 1 was not lower for those who selected peers who chose the Manual mode ($p = .408$). Model (2) reveals that Manual-Dishonest individuals were more likely to create homophilous links than those who chose Automatic were, and the Manual-Honest participants did not differ from the latter. Similar tests, performed separately on honest and on dishonest participants, reject homophilous link formation at the 5% level for Manual-Honest participants ($p = .079$, model (5)) but not for Manual-Dishonest participants ($p < .001$, model (6)).

Our analysis provides additional insight into the mechanism behind the observed homophilous link formation. If it was driven by a pure taste for similarity, we should observe it for both honest and dishonest individuals. The fact that it characterizes only individuals with a weaker moral type and neither participants who chose the Automatic mode nor those who chose the Manual mode but did not lie, suggests that the homophilous link formation was rather driven by strategic motives such as a self-serving willingness to bond with a reference group that signals a weak empirical social norm.²¹

Additionally, we can reject that the homophilous link formation was driven by social image concerns because the experimenter knew the lie size, by design, and the reported performance was not systematically below peers' performance and very rarely equal to it (see footnote 17). Finally, we reject that link formation was driven by the proximity with peers in terms of performance (instead of mode). This is clear for the players who chose the Automatic mode because higher proximity in performance would have led to the selection of peers who also chose the Automatic mode, which our analysis rejects. In addition, despite a potential endogeneity issue, we estimated whether the Manual-Honest individuals were less likely to select peers who chose the Manual mode when their own part 1 actual performance was closer to the mean performance of the peers who chose the Automatic mode. As can be seen from Table A.6 in Appendix, homophilous link formation either in terms of performance or in terms of mode is rejected for this sample, as in model (5) in Table 3.

5.4. Impact of homophilous link formation on the measure of peer effects

We now test Proposition 3 by checking the presence in ENDO of a self-selection bias because of the homophilous link formation we just observed, or of a link strength effect. If present, both the self-selection bias and the link strength effect would inflate the size of peer effects in this treatment, compared with those measured in EXO in subsection 5.2. Using a joint test on participants who chose the Manual mode and selected peers who made the same choice, we show that peer effects are of the same magnitude in EXO and ENDO. We conclude that homophilous link formation did not result in a bias and did not inflate our estimates of peer effects, which rejects Proposition 3.

Result 4. (Peer Effects, Selection Bias, and Link Strength) In the ENDO treatment, we jointly reject the presence of a selection bias and a link strength effect generated by the homophilous link formation in the measure of peer effects.

Support for Result 4. Following Eq. (20), we estimate a RE Tobit model on the pooled samples from the two treatments. We regress the size of lies on the average reported performance of peers in EXO and ENDO and add as regressors the average reported performance of peers in ENDO and a dummy variable (D_{ENDO}) for ENDO while controlling for the usual individual characteristics. Table 4 reports the results.

Model (1) in Table 4 confirms that the marginal effect of the performance of peers who chose the Manual mode is positive and highly significant. On average, in the pooled EXO and ENDO treatments an increase of 1 earnings point reported by peers in a period increased the lies of participants by 0.191 earnings points. The coefficient of the $(\text{MANUAL})\text{Peers' Reported Performance} \times D_{ENDO}$ variable is positive but not significant (and small). A one-tail test jointly rejects that the homophilous link formation generated a self-selection bias and a link strength effect that would affect the size of peer effects in ENDO but not in EXO. The dummy variable (D_{ENDO}) is not significant, indicating that ENDO had no effect *per se* on the size of lies. However, homophilous link formation is found only for the participants who chose the Manual mode and lied in part 1. Thus, we re-estimate model (1) separately for the Manual-Honest (model (2)) and the Manual-Dishonest participants (model (3)).²² Table 4 confirms Result 2 that peer effects are not significant for the Manual-Honest

²¹ Some participants mentioned in the ex-post questionnaire that they selected peers who chose the Manual mode as a source of motivation to perform better. However, this is only anecdotal, as only eight participants referred to this argument (six among those who chose the Automatic mode for themselves and two among those who chose the Manual mode). Moreover, this motive may refer more to the search of a reference point than to a willingness to compete (especially since there is no monetary incentives attached to overperforming peers). Finally, the performance of the peers who chose the Manual mode had an impact on the size of lies of individuals, but not on their actual performance. Thus, it is unlikely that competition is an important driver of behavior.

Table 4

Impact of endogenous matching on the estimated peer effects.

Dependent Variable: Size of lies	MANUAL-ALL (1)	MANUAL-H (2)	MANUAL-D (3)
(MANUAL) peers' reported performance	0.191*** (0.061)	0.008 (0.023)	0.201** (0.093)
(MANUAL) peers' reported performance $\times D_{ENDO}$	0.006 (0.088)	-0.018 (0.034)	0.068 (0.144)
D_{ENDO}	-1.444 (2.197)	-0.005 (0.670)	-4.728 (3.557)
Individual characteristics	Yes	Yes	Yes
N	730	345	385
Left-censored	292	275	17
Log-likelihood	-1487.10	-282.11	-1135.11

Notes: The table reports the average marginal effects over time and individuals from RE Tobit regressions on EXO and ENDO. Standard errors are in parentheses. The dependent variable is the size of lies in each period of part 2. There are five observations per individual. (MANUAL) Peers' reported performance is the average reported performance observed in the period from the two peers who chose Manual. D_{ENDO} is a dummy variable equal to 1 if the treatment is ENDO, and to 0 if it is EXO. The individual characteristics are the same as in the previous regression tables. Only the *First participation* variable is significant at the 1% level and negative in model (1). Model (1) includes all the participants who chose Manual; models (2) and (3) include only the participants who chose Manual and were honest in part 1 or, respectively, dishonest.

** $p < 0.05$ *** $p < 0.01$.

participants but significant and strong for the Manual-Dishonest participants; for the latter category, peer effects are equal in EXO and ENDO, confirming the finding of model (1).

6. Discussion and conclusion

In empirical studies, when using observational data the estimation of peer effects is often blurred because individuals may behave similarly to their peers partly because similar individuals tend to link together. In the extreme, peers could have the same behavior in the absence of any interindividual influence, simply because of homophilous matching. By contrast, a laboratory setting allows researchers to observe perfectly the formation of links and peer information. Through relevant treatment manipulations, it helps identify separately the effect of the endogenous matching and the effect of peers' behavior on decisions motivated by conformity. This is precisely one of the novelties of this study.

The results from our experiment suggest that i) causal peer effects affected the lying behavior of individuals who already lied in isolation; ii) homophilous link formation only manifested among liars; and iii) we jointly reject that such homophilous link formation biases the measure of conformity and that endogenous matching generates a link strength effect, as peer effects were not higher when individuals could select their peers. The absence of a severe self-selection bias generated by the homophilous link formation contrasts with some of the non-experimental papers that have suggested that endogeneity in network formation is a source of bias in the peer effect estimator (e.g., Carrell et al., 2013; Hsieh et al., 2020a). Our finding is consistent with those of other studies, in particular the ones using observational data from Add Health (e.g., Goldsmith-Pinkham and Imbens, 2013; Boucher, 2016; Badev, 2021), that found a weak effect of network endogeneity on this estimator. As long as there is no strategic complementarity, our analysis suggests that homophilous self-selection is not likely to be a serious source of bias on the estimation of conformity effects. The fact that we do not find assortative matching for both honest and dishonest individuals in our experiment does not affect this conclusion.

Our study delivers other interesting findings. In particular, not all individuals lie even after choosing an environment that enables them to misreport their outcome to earn more. The novelty is that this is found in a setting in which individuals can choose their cheating opportunities. This is suggestive of the presence of a strong social norm of honest reporting. Indeed, such a social norm could explain that honest individuals are unaffected by others' reported performance and by the treatment. Another notable result is that behavior was not influenced by the actual performance of peers who chose a setting that forbids lying. It was influenced by the reports of peers who chose a setting that allows lying, but this was the case only among individuals who lied when working in isolation. An interpretation is that individuals with a stronger moral type or a high willingness to comply with the dominant social norm of honesty are indifferent to the opportunities provided by the flexible environment when they are confident in their ability to surrender temptation. For some, choosing a setting

²² As a robustness test, we replicated Table 4 with a stricter definition of Manual-Dishonest participants. In Table A.7 in the Appendix, an individual is classified as Manual-Dishonest if he reported on average more than 2 more points than his actual performance per period in part 1 (instead of more than 1 point in Table 4), and as Manual-Honest otherwise. The results remain qualitatively similar.

that allows them to lie and not lie might be motivated by the willingness to signal their moral type. Consistently, when matching was exogenous, individuals who behaved honestly were not influenced by the reported performance of potential liars. Thus, it may not be surprising that their matching choices were not homophilous.

By contrast, the individuals who had revealed a weak moral type by lying in isolation were influenced by peers who chose the same mode. Consistently, a higher share of these individuals was willing to be matched with such peers. These different matching choices between dishonest and honest individuals suggest that the mechanism behind homophilous link formation was not a taste for similarity (which should be common across types of individuals), but rather a self-serving selection by liars of their source of information. Such link formation provides wilful ignorance of honest reports by avoiding repeat comparisons with peers' actual performance, and connects individuals with a reference group signalling a less virtuous norm. If the behavior of honest individuals is driven by the willingness to comply with a strong social norm of honesty, dishonest individuals might be all the more in need of observing empirical information that would contradict such honesty norm, which might contribute to explain why they selected similar peers.

Because information on peers' performance is also information on peers' payoffs, other potential mechanisms could be inequality aversion, status concerns, or competitiveness. However, there is no evidence that inequality aversion is what motivated lying because individuals who were more prone to peer effects and homophilous matching were those who already lied in isolation, because we found no peer effects on the actual performance, and because the size of lies rarely equalized the individual's and his peers' performance. Also, status concerns would have implied that individuals reacted similarly to peer performance in the Manual and Automatic modes and that individuals would be more likely to select peers who chose the Automatic mode since these peers would be easier to overcome. On the opposite, a taste for competition could lead individuals, even the honest ones, to select potential dishonest peers because they are more challenging; if this was the case, peers' performance would be expected to have an impact on the individuals' own performance. This is not what was observed in our data. Thus, we reject status seeking and competitiveness as the main drivers of our findings.

We also reject that our results are driven by social learning since neither the task nor lying required any sophistication. We acknowledge, nevertheless, that participants only received an imperfect signal about their peers' moral values and this may have influenced the size of peer effects through belief formation and updating. We gave an imperfect signal and we did not elicit beliefs on purpose to avoid making lying too salient and creating an experimenter demand effect. To measure the importance of this feature, a possible extension could be varying the degree of precision of the signal about peers' moral type before individuals create links and make decisions in a network.

In the field like in the laboratory, peer effects are manifest in the domain of morality but also in many other real-world settings, as diversified as addictive behavior, junk food consumption, or classroom achievements. Our experiment shows that peer effects increased lying through a contamination effect, but this affected only individuals who were already misbehaving in isolation and were probably in search of examples of social norm violations. Although we remain cautious about their external validity, these findings suggest that in more general real-world settings, it might be important to diversify the composition of peers networks such that those who are more inclined to norm violations are exposed to peers who are complying with the dominant social norm, such that they cannot avoid "good examples". We also suggest additional directions for future research. A notable extension would be to introduce incentive schemes in which payoffs depend on peers' performance. It would allow us to measure the extent to which conformity effects differ from those of strategic complementarities, and whether endogenous matching would affect these effects differently. In addition, we interpret the homophilous link formation by liars in terms of the choice of a reference group with a more lenient social norm, but we did not directly elicit the social norm. Further research could explicitly measure the impact of homophily on the perceived norm. Finally, testing how observability and communication would affect link formation and its impact on peer effects would open interesting perspectives for future research.

Declaration of Competing Interest

The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.jebo.2022.06.024](https://doi.org/10.1016/j.jebo.2022.06.024).

References

- Abeler, J., Nosenzo, D., Raymond, C., 2019. Preferences for truth-telling. *Econometrica* 87 (4), 1115–1153.
- Acemoglu, D., Ozdaglar, A., Siderius, J., 2021. Misinformation: strategic sharing, homophily, and endogenous echo chamber. NBER Working Paper 28884.
- Ajzenman, N., 2021. The power of example: corruption spurs corruption. *American Economic Journal: Applied Economics* 13 (2), 230–257.
- Akerlof, G.A., Kranton, R.E., 2000. Economics and identity. *Q J Econ* 115 (3), 715–753.
- Akin, Z., 2019. Dishonesty, social information, and sorting. *Journal of Behavioral and Experimental Economics* 80, 199–210.
- Aral, S., Walker, D., 2014. Tie strength, embeddedness, and social influence: a large-scale networked experiment. *Manage Sci* 60, 1352–1370. doi:[10.1287/mnsc.2014.1936](https://doi.org/10.1287/mnsc.2014.1936).

- Baccara, M., Yariv, L., 2013. Homophily in peer groups. *American Economic Journal: Microeconomics* 5 (3), 69–96.
- Badev, A., 2021. Nash equilibria on (un)stable networks. *Econometrica* 89 (3), 1178–1206.
- Bäker, A., Mechtel, M., 2019. The impact of peer presence on cheating. *Econ Inq* 57 (2), 792–812. doi:10.1111/ecin.12760.
- Ballester, C., Calvo-Armengol, A., Zenou, Y., 2006. Who's who in networks: wanted: the key player. *Econometrica* 74 (5), 1403–1417.
- Benhabib, J., Bisin, A., Jackson, M.O., 2010. The handbook of social economics, Vol. 1A and 1B. Amsterdam: North Holland.
- Bicchieri, C., Dimant, E., Gächter, S., Nosenzo, D., 2020. Social proximity and the erosion of norm compliance. IZA DP No. 13864.
- Bock, O., Baetge, I., Nicklisch, A., 2014. Hroot: hamburg registration and organization online tool. *Eur Econ Rev* 71, 117–120.
- Bott, K.M., Cappelletti, A.W., Sorensen, E., Tungodden, B., 2020. You've got mail: a randomized field experiment on tax evasion. *Manage Sci* 66 (7), 2801–2819.
- Boucher, V., 2016. Conformism and self-selection in social networks. *J Public Econ* 136, 30–44.
- Boucher, V., Fortin, B., 2016. Some Challenges in the Empirics of the Effects of Networks. In: Bramoullé, Y., Galeotti, A., Rogers, B. (Eds.), *The Oxford Handbook of the Economics of Networks*. Oxford University Press, pp. 277–302.
- Bramoullé, Y., Currarini, S., Jackson, M.O., Pin, P., Rogers, B.W., 2012. Homophily and long-Run integration in social networks. *J Econ Theory* 147 (5), 1754–1786.
- Calvo-Armengol, A., Patacchini, E., Zenou, Y., 2009. Peer effects and social networks in education. *Rev Econ Stud* 76 (4), 1239–1267.
- Carrell, S.E., Malmstrom, F.V., West, J.E., 2008. Peer effects in academic cheating. *Journal of Human Resources* 43 (1), 173–207.
- Carrell, S.E., Sacerdote, B.I., West, J.E., 2013. From natural variation to optimal policy? the importance of endogenous peer group formation. *Econometrica* 81 (3), 855–882.
- Currarini, S., Jackson, M.O., Pin, P., 2009. An economic model of friendship: homophily, minorities, and segregation. *Econometrica* 77 (4), 1003–1045.
- Currarini, S., Mengel, F., 2016. Identity, homophily and in-group bias. *Eur Econ Rev* 90, 40–55.
- Damm, A., Dustmann, C., 2014. Does growing up in a high crime neighborhood affect youth criminal behavior? *American Economic Review* 104 (6), 1806–1832.
- Dimant, E., 2019. Contagion of pro- and anti-social behavior among peers and the role of social proximity. *J Econ Psychol* 73, 66–88.
- Drago, F., Mengel, F., Traxler, C., 2020. Compliance behavior in networks: evidence from a field experiment. *American Economic Journal: Applied Economics* 12 (2), 96–133.
- Dufwenberg, M., Dufwenberg, M.A., 2018. Lies in disguise a theoretical analysis of cheating. *J Econ Theory* 175, 248–264.
- Daz, C., Patacchini, E., 2020. Parents, neighbors and youth crime. IZA DP No. 13906.
- Efferson, C., Lalive, R., Fehr, E., 2008. The coevolution of cultural groups and ingroup favoritism. *Science* 321 (5897), 1844–1849.
- Egan, M., Matvos, G., Seru, A., 2019. The market for financial adviser misconduct. *Journal of Political Economy* 127 (1), 233–295.
- Falk, A., Huffman, D., Mierendorff, K., 2006. Incentive effects and political acceptability of workfare. Institute for the Study of Labour (IZA).
- Fischbacher, U., Föllmi-Heusi, F., 2013. Lies in disguise an experimental study on cheating. *J Eur Econ Assoc* 11 (3), 525–547.
- Fiske, S.T., 2018. Social beings: Core motives in social psychology. John Wiley & Sons.
- Flashman, J., Gambetta, D., 2014. Thick as thieves: homophily and trust among deviants. *Rationality and Society* 26 (1), 3–45.
- Fortin, B., Lacroix, G., Villeval, M.C., 2007. Tax evasion and social interactions. *J Public Econ* 91 (11), 2089–2112.
- Fosgaard, T.R., Hansen, L.G., Piovesan, M., 2013. Separating will from grace: an experiment on conformity and awareness in cheating. *Journal of Economic Behavior & Organization* 93, 279–284.
- Gavrilova, E., 2019. A partner in crime: assortative matching and bias in the crime market. *Journal of Economic Behavior & Organization* 159(C), 598–612.
- Gino, F., Ayal, S., Ariely, D., 2009. Contagion and differentiation in unethical behavior: the effect of one bad apple on the barrel. *Psychol Sci* 20 (3), 393–398.
- Glaeser, E.L., Sacerdote, B., Scheinkman, J.A., 1996. Crime and social interactions. *Q J Econ* 111 (2), 507–548.
- Gneezy, U., Kajackaite, A., Sobel, J., 2018. Lying aversion and the size of the lie. *American Economic Review* 108 (2), 419–453.
- Goldsmith-Pinkham, P., Imbens, G.W., 2013. Social networks and the identification of peer effects. *Journal of Business & Economic Statistics* 31 (3), 253–264.
- Golub, B., Jackson, M.O., 2012. How homophily affects the speed of learning and best-response dynamics. *Q J Econ* 127 (3), 1287–1338.
- Graham, B.S., 2019. Network data. Technical Report. National Bureau of Economic Research.
- Gross, J., Leib, M., Offerman, T., Shalvi, S., 2018. Ethical free riding: when honest people find dishonest partners. *Psychol Sci* 1956–1968.
- Haynie, D.L., 2001. Delinquent peers revisited: does network structure matter? *American Journal of Sociology* 106 (4), 1013–1057.
- Henrich, J., Boyd, R., 1998. The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior* 19 (4), 215–241.
- Hsieh, C.-S., Knig, M.D., Liu, X., Zimmermann, C., 2020. Collaboration in bipartite networks, with an application to coauthorship networks. Technical Report. Working Paper 2020-030A, Federal Reserve Bank of St-Louis.
- Hsieh, C.-S., Lee, L.-F., Boucher, V., 2020. Specification and estimation of network formation and network interaction models with the exponential probability distribution. *Quant Econom* 1349–1390.
- Innes, R., Mitra, A., 2013. Is dishonesty contagious? *Econ Inq* 51 (1), 722–734.
- Irlenbusch, B., Mussweiler, T., Saxler, D.J., Shalvi, S., Weiss, A., 2020. Similarity increases collaborative cheating. *Journal of Economic Behavior & Organization* 178, 148–173.
- Keizer, K., Lindenberg, S., Steg, L., 2008. The spreading of disorder. *Science* 322 (5908), 1681–1685.
- Kiessling, L., Radbruch, J., Schaub, J., 2019. Self-selection of peers and performance. Discussion Paper University of Bonn and University of Mannheim 121.
- Knecht, A., Snijders, T., Baerveldt, C., xxx, 2010. Friendship and delinquency: selection and influence processes in early adolescence. *Social Development* 19, 494–514.
- Kocher, M.G., Schudy, S., Spantig, L., 2018. I lie? we lie! why? experimental evidence on a dishonesty shift in groups. *Manage Sci* 64 (9), 3995–4008.
- Konrad, K.A., Lohse, T., Simon, S.A., 2021. Pecunia non olet: on the selfselection into (dis)honest earning opportunities. *Experimental Economics* Forthcoming.
- Kroher, M., Wolbring, T., 2015. Social control, social learning, and cheating: evidence from lab and online experiments on dishonesty. *Soc Sci Res* 53, 311–324. doi:10.1016/j.ssresearch.2015.06.003.
- Lauer, T., Untertrifaller, A., 2019. Conditional dishonesty. Mimeo doi:10.13140/RG.2.2.30604.16006.
- Lazear, E., Malmendier, U., Weber, R., 2012. Sorting in experiments with application to social preferences. *American Economic Journal: Applied Economics* 4 (1), 136–163.
- Lefebvre, M., Pestieau, P., Riedl, A., Villeval, M.C., 2015. Tax evasion and social information: an experiment in Belgium, France, and the Netherlands. *International Tax and Public Finance* 22 (3), 401–425.
- Leib, M., Schweitzer, M., 2020. Peer behavior profoundly influences dishonesty: will individuals seek-out information about peers dishonesty? Mimeo.
- Li, S.X., 2020. Group Identity, Ingroup Favoritism, and Discrimination. In: Zimmermann, K. (Ed.), *Handbook of Labor, Human Resources and Population Economics*. Springer.
- Manski, C.F., 1993. Identification of endogenous social effects: the reflection problem. *Rev Econ Stud* 60 (3), 531–542.
- McPherson, M., Smith-Lovin, L., Cook, J.M., 2001. Birds of a feather: homophily in social networks. *Annu Rev Sociol* 27 (1), 415–444.
- Page, T., Putterman, L., Unel, B., 2005. Voluntary association in public goods experiments: reciprocity, mimicry and efficiency. *The Economic Journal* 115 (506), 1032–1053.
- Pascual-Ezama, D., Dunfield, D., Gil-Gomez de Liano, B., Prelec, D., 2015. Peer effects in unethical behavior: standing or reputation? *PLoS ONE* 10 (4), e0122305.
- Patacchini, E., Rainone, E., Zenou, Y., 2017. Heterogeneous peer effects in education. *Journal of Economic Behavior & Organization* 134, 190–227.
- Patacchini, E., Zenou, Y., 2009. Juvenile delinquency and conformism. *Journal of Law, Economics, & Organization* 28 (1), 1–31.

- de Paula, A., 2017. Econometrics of Network Models. In: Honoré, B., Pakes, A., Piazzesi, M., Samuelson, L. (Eds.), *Advances in Economics and Econometrics: Eleventh World Congress*, Vol. 1. Cambridge University Press, pp. 268–323.
- Rauhut, H., 2013. Beliefs about lying and spreading of dishonesty: undetected lies and their constructive and destructive social dynamics in dice experiments. *PLoS ONE* 8 (11), e77878.
- Ushchev, P., Zenou, Y., 2020. Social norms in networks. *J Econ Theory* 185, 104969.
- Xiang, R., Neville, J., Rogati, M., 2010. Modeling relationship strength in online social networks. In: *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, pp. 981–990. doi:[10.1145/1772690.1772790](https://doi.org/10.1145/1772690.1772790).