Experimental Evidence on the Transmission of Honesty and Dishonesty: A Stairway to Heaven and a Highway to Hell

Paola Colzani	**Georgia Michailidou†**	**Luis Santos-Pinto**
University of Lausanne	New York University Abu Dhabi	University of Lausanne
Switzerland	*United Arab Emirates*	*Switzerland*

September 2022

Abstract

Theories of social behavior propose that individuals condition actions that involve a moral value by following others' behavior. The theoretical and experimental instruments evaluating this conditioning often focus on actions with negative moral value (e.g. dishonesty, norm violation, tax evasion). Here, we execute a laboratory experiment to evaluate the diffusion of actions with positive and negative moral values. We use a lying dilemma and introduce a novel methodology operationalizing beliefs as intention proxies to study the switch between honesty and dishonesty in simultaneous and sequential move sequences. The results indicate asymmetries; while lying is contagious, truth-telling is not.

JEL classification: C90, C91, C92, D90

Keywords: lying, truth-telling, moral behaviour, conditional behaviour, laboratory experiment, die-rolling paradigm

†Corresponding author: georgina@nyu.edu, ORCID: 0000000255450557

1. Introduction

Humans have developed various adaptation mechanisms to facilitate survival and procreation. One such mechanism is imitation. For many of our social actions we tend to follow the example set by others. When actions are observable, certain behaviours can pass on from one individual to the next and become norms. In this paper, we employ a laboratory experiment to study whether people follow others' example in matters of moral behaviour. In particular, we focus on the aspect of morality pertaining to honesty to consider whether both, lying and truth-telling, are transmittable among individuals. Evaluating the "contagiousness" of honesty and dishonesty is arguably an interesting endeavour in itself but it is also important for policymaking. If people imitate honest behaviour, then, setting the right moral examples in key positions could work as ethics diffusers. Similarly, if they also imitate dishonesty, building social architectures that hinder observability among potentially dishonest actors could help with dishonesty containment. Our results reveal an asymmetry; individuals imitate lying but not truth-telling.

The line of research we pursue builds upon the literature on lying[1] and truth telling. An earlier strand of the literature demonstrated that, contrary to standard economic

[1] Lying, in this paper, will refer to knowingly misreporting a clear and specific outcome. As we make no assumptions or reference to theory of mind, lying in our context could have been specified also as cheating, or

intuition, individuals exhibit preferences for honesty. Even in situations in which lying increases payoffs and does not produce pecuniary externalities[2], a non-negligible proportion would remain honest or would lie, but not fully (Fischbacher and Föllmi-Heusi, 2013; Abeler, et al., 2014). One of the explanations put forward is that honesty, as a social norm, has been internalized, making individuals lying averse due to the psychological costs accruing when one deviates from the social norm of honesty (Hurkens and Kartik, 2009; López-Pérez and Spiegelman, 2013; Kartik et al., 2014). Further, that individuals might engage in some partial lying to the extent that they can justify this lying and maintain a positive self-view of an honest person (Mazar et al., 2008; Shalvi et al.,2011; Shalvi et al., 2012). Recent evidence suggests that these individual preferences might not be for honesty per se, but rather, for appearing honest. Following, that individuals entertain social image concerns, which allow for a more flexible view of honesty. In contexts where lying cannot be precisely estimated, or liars cannot be identified, individuals are not as honest as models of pure preferences for honesty would predict (Gneezy et al., 2018; Abeler, et al., 2019).

In this paper, we examine whether honesty and dishonesty, in environments without payoff interdependencies, pass on from one individual to another. That is, whether

deceiving. However, these terms carry their own definitions and, in some contexts, they have distinct properties. For a refined and thorough categorization of such definitions, see Sobel, 2020.

[2] When lying produces negative pecuniary externalities for others, a large fraction of individuals also refrain from lying (Gneezy, 2005)

individuals condition their lying behaviour on that of others, in situations where the benefits from lying or truth-telling are independent but observable among individuals. The notion of conditional morality, i.e. that people imitate others' lying and truth-telling behaviour, is consistent with the theory of social image concerns and with the standard economic assumption of self-interest[3]. Although it departs from the concept of honesty as an internalized norm, it does imply the existence of both normative and empirical expectations, which might not always overlap[4]. In a plain storyline, conditional morality works as follows; individuals are self-interested and have the conflicting desire to obtain the higher payoffs produced from lying, but also, to maintain the social and self-image of an honest person. For low payoffs, social and self-image concerns are dominating payoff maximizing motivations and individuals behave honestly, as they *ought to* (Kajackaite and Gneezy 2017). This is true unless individuals expect or know that others are lying. These expectations, or knowledge, can shift focus from what *ought to be done* to what *is being done* by others and draw a moral wiggle room for individuals to behave dishonestly without suffering social image discomforts or shame and guilt. Similar arguments can apply on the reverse. In cases where individuals expect or know that others respond to lying dilemmas with truth-telling, and own actions are observable by those others, social-

[3] To this regard, social conformity on moral issues and conditional morality overlap as definitions in our context. We chose to go with the term conditional morality because it highlights the moral component of the environment we are considering.
[4] These expectations might also be strategically distorted by oneself in order to accommodate a conditional reasoning of one's lying (Bicchieri et al., 2020). That is, individuals might make themselves believe that lying is more widespread to provide themselves an excuse to lie in subsequent lying opportunity.

image concerns, shame, and guilt become more salient and reinforce what *ought to be done*. Now, there is no moral wiggle room; rather, an audience of honest others who might silently pass moral judgment since they themselves could have lied but did not.

To evaluate whether conditional morality exists in the absence of payoff interdependencies, we execute an experiment which features a between-subjects design with two treatments. In one treatment, individuals are in pairs and play a simultaneous move lying dilemma. In the second treatment, individuals, again in pairs, play a sequential move lying dilemma. This allows the second mover to obtain knowledge about the moral behaviour of the first mover. If individuals are prone to follow others' example, we should find traces thereof in two ways (i) The lying behaviour of all subjects of the simultaneous treatment and the first movers of the sequential treatment should align with their beliefs about peers' lying behaviour. This means that those who respond honestly should expected others to respond honestly and those who respond dishonestly should expected others to respond dishonestly. (ii) The second movers of the sequential treatment should follow the lying behaviour of the first movers they are paired with. That is, those who observe first movers lying should lie, and those who observe first movers being honest should be honest irrespectively of whether their beliefs about the first movers' behaviour were confirmed or not.

Our findings show that conditional morality is asymmetric; while lying is contagious, truth-telling is not. Indeed, beliefs and behaviour correlate significantly, meaning that honest subjects believe others are honest and dishonest ones believe others are dishonest. However, when beliefs become falsified, behavioural conditioning applies selectively. Subjects who expected others to tell the truth and whose expectations were falsified, lie significantly more than those who expected others to tell the truth and whose expectations were confirmed. In contrast, subjects who expected others to lie and whose expectations were falsified, lie as much as those who expected others to lie and whose expectations were confirmed.

2. Contribution and Related Literature

Conditional morality, either as act of imitation, or as peer effects, or as social learning has been studied extensively, theoretically and experimentally[5]. Our study contributes to the literature on the behavioral conditioning of lying and truth-telling. Moreover, it makes two important methodological contributions to this literature. First, by considering both positive and negative conditioning, and second, by capturing the switching from intended to exhibited behavior post conditioning. More on the former, we separate the conditioning effect by the moral connotation of the preceding act, meaning the focus is not only on those who anticipate or observe anti-social behavior, but also on those who anticipate or observe pro-social behavior. It is important to separate and consider this

[5] The related literature is presented and discussed thoroughly in the following subsections.

duality. Often in the literature, behavior is studied on an aggregate level without considering the positive spillovers that might be stemming from honest conduct. This can lead to under-estimation of the conditioning effect. Given our design, we can identify the connotation of both the preceding and the consequent act, meaning, we can observe lying and truth-telling on an individual level.

Since we are able to identify the interplay between anticipated and observed actions, we can evaluate the likelihood with which individuals follow the moral example of the observed act and whether this likelihood differs when the observed act contradicts either an honest or a dishonest *anticipated* act. In other words, this study's distinctive contribution is that it employs a methodology that allows us to study behavioral switching. By recording the exhibited differences between (honest) individuals who expected others to be honest but observed others being dishonest, with (dishonest) individuals who expected others to be dishonest but observed others being honest.

Below we discuss the related literature. As this is an experiment on honesty and dishonesty, we start with experimental evidence on lying from the laboratory. Closest to ours, are lying experiments in which subjects receive information about others' behavior. Although such experiments exist, they either involve strategic interactions which confound moral diffusion streams, or (dis)honesty is not observable on an individual level making impossible to identify behavioral switching, or they only consider the diffusion of dishonesty without considering the spillovers from truth telling. We also discuss studies conceptually related to ours but that focus on prosocial and antisocial behaviors in different contexts, i.e., tax compliance, public goods, charitable giving, and cooperation.

2.1. Literature on Lying behavior in the laboratory

In the lab, lying behavior has been studied extensively using the die-roll paradigm introduced by Fischbacher and Follmi-Heusi, 2013. In their design, subjects privately observe the outcome of a die-roll and are asked to make a report that will determine individual monetary payoffs. Fischbacher and Follmi-Heusi, 2013, as well as Gachter and Schulz, 2016, and Gneezy et al., 2018 find that people lie but not maximally. Lying behavior has also been investigated under different types of strategic interactions, such as deception games and coordination games. In deception games, subjects strategically choose to lie taking into consideration not only their own gains but also the possibility to favor or damage others with their choice (Gneezy, 2005; Erat and Gneezy, 2012; Gneezy et al., 2013). In coordination games, people strategically lie more if the monetary benefits of lying are equally shared with others i.e. there is a payoff commonality (Conrads et al., 2014; Weisel and Shalvi, 2015; Barr and Michailidou, 2017) or when they can communicate (Kocher, Schudy and Spantig 2016). In coordination games where the benefits of lying are not equally shared, Lauer and Untertrifaller, 2019, find that one third of the subjects engage in strategic lying if and only if one or more group members lie. In a related, lying coordination game in which three group members need to match each other's reports, Rilke et al., 2021, report that the presence of a first mover decreases dishonesty in repeated interactions but not in one-shot settings, a finding the authors attribute to the first movers' image concerns. Our study differs from Rilke et al., 2021, in that it involves a non-strategic decision-making context but, most importantly, a context in which observability and image concerns, are held constant across simultaneous and sequential treatments and thus cannot explain our findings. The study closest to ours is Diekmann et al, 2015. They study the impact of observing lying in the absence of strategic interactions by asking subjects to report twice, once without information and a second time after observing the distribution of reports in the first round of the experiment or after observing the distribution of reports from a similar Fischbacher and Follmi-Heusi, 2013, die-rolling experiment. The results show that observing the distribution of reports either in the first round or in a similar experiment increases lying. We are interested in understanding what determines lying behavior in absence of strategic lying confounds and when group members can precisely observe each other's lies. Such a framework allows us to study whether observing a group member lying increases lying and isolate reputational cost of lying from other lying motives. Further, with our design, we are able to exclude strategic interactions but also, we allow our subjects to know precisely whether and to what extent their group peers lied. Our results are in line with Diekmann et al., 2015 but extend them; observing lying increases lying but observing truth-telling does not increase truth-telling.

2.2. Literature on social conformity and tax compliance

Answering a similar research question Lefebvre et al., 2015 study the impact of positive and negative examples of peers' tax compliance behavior on the likelihood to tax evade. In line with our findings, they report that examples of high compliance do not produce a disciplining effect, but examples of low compliance significantly increase tax evasion. Also on tax evasion and social interactions, Fortin and Villeval 2007, report no evidence of social conformity when considering the impact of information on tax evasion on aggregated behavior.

2.3. Literature on peer effects, public goods provision, and charitable giving

Public good contributions and charitable giving are arguably different from honesty. Among other differences, contributions entail engaging in an action (giving) while honesty entails refraining from an action (lying). Although this difference does not reflect precisely a commission vs omission set up[6], there are still interesting comparisons to be drawn as the effort exerted in making contributions might give rise to different behavioral conditioning mechanisms than the effort from refraining the temptation to lie.

This is partly true considering the findings reported in Croson and Shang, 2008, and Shang and Croson, 2009. Both studies consider the impact of social information on contributions to public goods and use the same field experiment to do so. Both suggest

[6] Typically, omission is morally charged, but here, honesty is the morally virtuous stance.

that social information (i.e., information about peers' contributions) significantly affects individuals' behavior. When members of a public radio station renewed their annual contribution, they gave more if they learned other members gave more, or gave less when other members gave less. Compared to this study, we see that the asymmetries that we report, are not present in this field experiment. That is, observing others' higher contributions increases own contributions, whilst in our experiment, observing others' honesty does not drive own honesty. However, one needs to note that a reduced contribution is still a contribution (i.e. a pro social action) and thus cannot be directly compared to our set up in which lying comes with a negative moral connotation.

In another lab experiment, individuals also exhibit signs of contingent contribution, this time by following a leader's contribution. In this public goods game, Gächter and Renner, 2018, show how a first contributor (leader) initiates a path dependency of followers' contributions. Observing the leader's initial contribution, followers form beliefs about peers' contribution levels which in turn affect own contribution choices. Clusters with prosocial leaders sustain high contribution levels, while less prosocial leaders compel consistently low contributions that remain rigid even if leaders increase subsequent contributions.

In line with the asymmetries we report are the results from a gift exchange laboratory experiment by Thöni and Gächter, 2015. In their study, subjects can choose to revise their effort provision after learning the effort level chosen by their peers. Overlapping with our asymmetries, the authors report that subjects revise their effort downwards when their peers had chosen lower effort levels than themselves, however, they do not revise their effort upwards when their peers had chosen a higher effort level than themselves. We see that this “convenient conditioning” in which individuals will only follow social information when it is favorable to do so manifests itself in studies with moral implications (such is ours) as well as in studies without moral implications but costly personal choices.

2.4. Literature on social misbehavior begetting social misbehavior

Focusing on the contagiousness of antisocial behavior, a rich literature studies conditional social norm violation i.e. the higher propensity of individuals to break a rule when they observe others’ misbehavior. Keizer et al., 2008 find empirical evidence that being aware of others misconduct triggers more misconduct, for instance, observing illegal parking causes illegal trespassing, while Cialdini et al., 1990 show that people are more prone to litter when they observe others littering. Gino et al., 2009 find that people are more likely to misbehave when they observe an in-group breaking a rule but are more likely to adhere to such a rule when an out-group misbehaves. In a more recent study, Smerdon et al., 2020, show that pluralistic ignorance and the strength of social interactions create favorable conditions for “bad” social norms to persist. This is a particularly alarming finding since it suggests that individuals do not necessarily need to observe antisocial contact to behave antisocially. Rather, the pluralistic ignorance channel proposes that individuals only need to maintain wrong beliefs about what the prevailing norm is. If individuals believe bad norms are prevalent and are uncertain about others preferences, then these norms can persist and solidify.

Finally, Hugh-Jones and Ooi, 2017, show that norm diffusion starts early on in individuals' lives. Children and teenagers affect each other's notion of fairness through social interactions. Redistributive preferences transmit and get adopted by observing young peers while the strength of the influence is subject to the social distance and friendship ties among them.

3. Experimental Design

We employ a variation of the standard die-rolling paradigm by Fischbacher and Föllmi-Heusi, 2013, in which we manipulate the sequence of reporting. It is a computerized, one-shot, between-subjects design, consisting of two treatments; a simultaneous move treatment and a sequential move treatment. Instructions are common knowledge and full anonymity protocol is applied across all conditions. For both treatments, subjects are randomly matched into groups of two, and each of the two subjects in a pair observes the same electronic die roll. In the simultaneous treatment, the two subjects in each pair report independently but simultaneously. In the sequential treatment, one subject reports first, while the other reports after observing her pair's report. In both treatments, we elicit subjects' beliefs about their pair's behavior in an incentivized manner.

The first thing to note about our design is that subjects in each pair observe the same die roll outcome. This serves the purpose of minimizing the effect of potential inequity considerations that might accrue to individuals observing different outcomes.

The second thing to note is that lying is verifiable, meaning that electronic die-roll outcomes are observed and recorded by the experimenter. Considering the findings of Gneezy et al., 2018, which suggest that the proportion of maximal lying is higher than that of partial liars when outcomes are verifiable, we expect an analogy of more maximal liars and fewer partial liars across our experimental conditions. Because this is a common feature across all the experimental conditions, we do not think it compromises in any way our treatment comparisons. To verify that the observability from the experimenter does not produce effects different from those described in the literature, we also design and execute a control treatment, which is comparable with our two baseline treatments and with Gneezy et. al., 2018, experiment. We find that, indeed, the results from our control treatment are in line with the literature and we provide the design of the control together the comparing analysis with the results of Gneezy et. al., 2018, in the appendix.

The third thing to note is that in the sequential treatment, the first movers' report is observable by the experimenter but also by the second mover of the pair. We thought this might give rise to a different type of audience effects, or make the existing audience effects more salient for first movers, which could affect their behavior. Aiming to maintain consistency across treatments and good controls, we build our experimental design around this feature. As a result, we made all subjects' reports observable by their pair, in both treatments[7]. Making everyone's reports known to their pair also helps eliminating any trust issues our subjects' might have entertained concerning the belief elicitation task. All experimental instructions are provided in the Appendix.

3.1. Simultaneous Treatment

Subjects enter the lab and are seated in individual, panel divided experimental booths. They are randomly paired, and assigned the role of either A or B in a random fashion. Both subjects in each pair observe the same die roll outcome in their computer screens. Their main task is to report the die roll outcome they observed. Each subject's payoff depends on their own–and only on their own- report. That is, there is no payoff interdependency between subjects in each pair. A report of 6 pays CHF 6, a report of 5 pays CHF5, a report of 4 pays CHF 4 and so on, regardless of the actual die roll outcome or the report of the other subject in the pair.

---

[7] We also keep this feature present in the control treatment.

Subjects also guess their pair's report. A correct guess pays 1 CHF. After subjects have entered their guesses, they are informed about each other's report. Finally, subjects answer a short, non-incentivized questionnaire on sociodemographic characteristics.

3.2. Sequential treatment

In this treatment, all conditions are identical to the simultaneous treatment apart from the sequence of reporting and disclosure of subjects' reports to each other. In particular, after a pair observes the common die roll, A goes first to guess B's report and report while B is waiting. Then, B is asked to guess A's report. Afterwards, B learns A's report (i.e., whether her guess was correct) and is asked to proceed with her own report. Then, B's report is revealed to A.

An important concern needs to be addressed at this point. While all subjects of the simultaneous treatment and As in the sequential treatment make their reports without knowing the success of their guess, Bs in the sequential treatment make their report after learning As report and thus, whether their guess was accurate. While this is a fundamental part of our design as it helps us document behavioral switches, one could argue that, since correct guesses pay 1 CHF, those Bs who guessed wrongly might be more willing to lie to compensate for the loss of 1 CHF. We have good ex-ante and ex-post reasons to believe this did not compromise our design. Ex-ante, we purposefully kept the payment for correct guesses quite low to attenuate the incentives for loss compensation. Ex-post, our results suggest that inaccurate guesses are not affecting lying. As a design check exercise, we compare average lying in the sequential treatment among Bs who guessed wrongly (1.74) and those who guessed correctly (1.14) and a two-sided t-test rejects that the difference is significant (p=0.11).

3.3. Procedures and Payments

We conducted all experimental sessions at the behavioral laboratory at University of Lausanne (LABEX). Subjects were recruited using ORSEE (Greiner, 2015) and the experiment was programmed in zTree (Fischbacher, 2007). Instructions were presented in oral and in written form, in French language. Subjects' understanding of the experimental task was tested with a series of questions prior to the task, and their demographic characteristics were collected with a post-task questionnaire. Treatments were executed in random order, although we started with a session of simultaneous treatment to serve as the "Experiment 1" pool of reports that would be used for the control session[8]. In total, 266 subjects participated in our experiment. 100 subjects participated in the simultaneous treatment; half of them were assigned to the role of A and the other half to the role of B, 55% of all were males[9], and the mean age was 21.6 years. 166 subjects participated in the sequential treatment; half of them were assigned to

[8] 138 subjects took part in the control treatment but only half of them were active. In the analysis we present in the appendix, we only consider the 69 active subjects of this control.

[9] Here, we refer to the self-identified gender of subjects.

the role of A and the other half to the role of B, 44.57% of them were males and mean age was 21.7 years [10].

Sessions lasted on average 30 minutes. Each subject received 10 CHF as a show up fee, and total earnings varied between 11 and 17 CHF. Payments were carried via pay-sheets which subjects carried outside the lab and were cashed by an experimenter who was not present during the experiment and could not identify subjects' experimental behavior. Payment receipts could not be linked to subjects' experimental behavior. This was a deception free experiment and full anonymity applied throughout all sessions.

4. Results

The main question we explore in this paper is whether honesty and dishonesty transmit among individuals and whether they transmit similarly. Here, we do not define transmission solely as the act of following a first mover's moral choice but rather, as the act of diverging from the moral choice one intended to pursue, and switching to the moral choice of a first mover. To establish this switching, we will instrument the relationship between beliefs and choices of subjects in the simultaneous treatment. In particular, we will evaluate whether and to what extent, beliefs and choices overlap. If, according to the hypothesis we formulate and in line with much of the literature, in the simultaneous

[10] The gender ratios are marginally different between treatments; a test of proportions yields a p=0.099.

treatment, those who expect others to lie lie, and those who expect others to be honest are honest, then, we have a good proxy for Bs intended moral choice in the sequential treatment. If we can infer Bs intended moral choices via their beliefs, and document their actual choice after they observe As choices, we will be able to evaluate whether switching occurs. For this, we will hypothesize that, in the sequential treatment, Bs who intended to be honest but observed a dishonest A, will act dishonestly, and those Bs who intended to be dishonest but observed an honest A, will act honestly.

Although this is the main path we will use to study whether honesty or dishonesty transmit, we will also evaluate the transmission via more direct channels, meaning we will also examine the extent to which Bs in the sequential treatment follow As example, without considering beliefs. In addition, we use the appendix to present a view of the aggregated results, such as comparative analysis of lying in the intensive and extensive margin across treatments, including the control. We also include a comparative analysis between our aggregate results and the results of experiments who employ similar, observability of lying paradigms. We start the analysis by stating our first hypothesis.

Hypothesis 1A*: In the simultaneous treatment, those who lie believed their pair lied, and those who told the truth believed their pair told the truth.*

To evaluate the correlation between behavior and beliefs we will assume that, if no correlation exists, both honest and dishonest subjects' beliefs are random. That is, we will assume that 50% of honest subjects believe their pair was honest, and 50% of dishonest subjects believe their pair lied. We will test these null hypotheses against the one-sided alternatives that significantly more than 50% of honest subjects believe their pair was honest and that significantly more than 50% of dishonest subjects believe their pair lied. In our simultaneous treatment, 71.4% of honest subjects believed their pair was honest, and 90.9% of dishonest subjects believed their pair was dishonest. Both these percentages are significantly higher than 50%, both with $p< 0.01$[11]. Given this result, we conclude that our hypothesis that beliefs and behavior correlate, is confirmed[12]. To substantiate further this claim, we also examine lying levels among subjects with diverse beliefs and test the following related hypothesis:

Hypothesis 1B*: In the simultaneous treatment, average lying of those who believed their pair told the truth is significantly lower than average lying of those who believed their pair lied.*

A two-sided t-test confirms this hypothesis. Average lying among those who expect others to tell the truth is 0.27 and it is significantly lower ($p<0.001$) from 2.05 which is the average lying of those who expect others to lie. Taken together, these two results allow us to use beliefs as a proxy for the intended actions of Bs in the sequential treatment. This means, we consider valid to assume that those Bs in the sequential treatment who expect

[11] Results remain the same if we consider two-sided tests.

[12] Let us highlight that no claims about causality are made in this part of our analysis. Behavior might be driving beliefs or beliefs might be driving behavior. For our subsequent analysis, it is necessary for us to establish correlation between beliefs and behaviour, not causality. A more elaborate analysis on how beliefs are becoming and sustained see Bicchieri et al., 2020.

others to be honest/dishonest, intended to act honestly/dishonestly, and we proceed to examine whether they switch their intended action if they observe an A who acted opposite to their expectation.

Hypothesis 2A –Transmission of Dishonesty*: In the sequential treatment, Bs who anticipated As to tell the truth and observed As who lied, lie significantly more than those who anticipated As to tell the truth and observed As who told the truth.*

Hypothesis 2B -Transmission of Honesty*: In the sequential treatment, Bs who anticipated As to lie and observed As who told the truth, lie significantly less than those who anticipated As to lie and observed As who lied.*

Considering hypotheses 2A and 2B together, we identify four categories of subjects: Bs who expected an honest A and A was honest, Bs who expected an honest A and A was dishonest, Bs who expected a dishonest A and A was dishonest, and finally, Bs who expected a dishonest A and A was honest. We depict these four categories in Figure 1 below.

In Figure 1 the bars stand for average lying and the spikes represent 95% confidence intervals. The results are quite revealing. Using the first two bars, we can infer what is the impact of observing dishonest As on the behavior of Bs who were expecting As to act honestly.

Bs who anticipated A to tell the truth and whose expectations were confirmed, lied on average 0.05 (first bar), but Bs who anticipated A to tell the truth and whose expectations were not confirmed, lied on average 2.05 (second bar). This is a strongly significant difference (p<0.001).

Figure 1: Average lying across anticipation and observation categories - Sequential Treatment

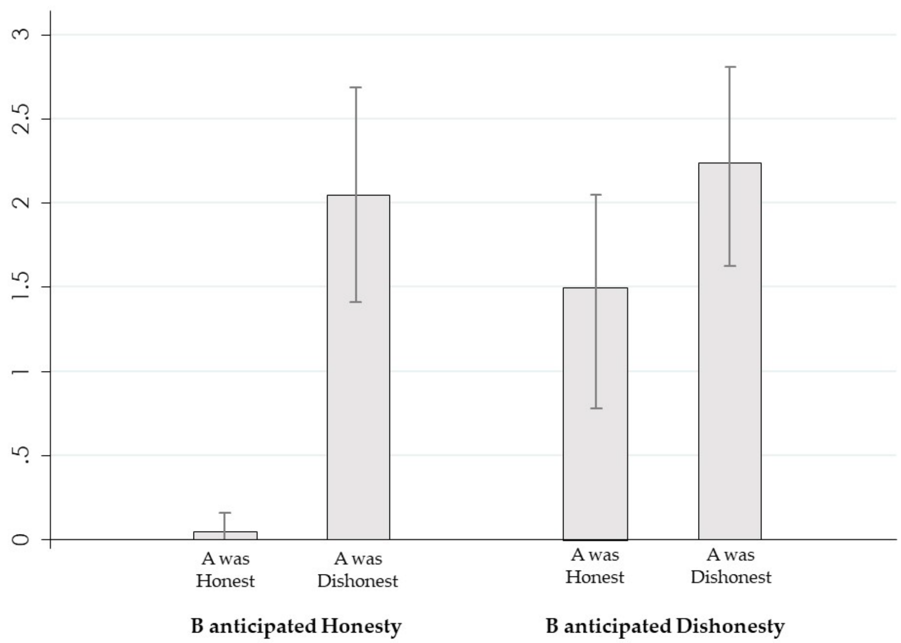

Guided by the simultaneous treatment, we would predict that those who believe others to be honest lie very little, and somewhere in the level of 0.27. This holds in the cases where anticipations about honesty are confirmed, however, when anticipations are disproved, Bs are ready to follow the example of dishonesty set by As and to lie to significantly higher levels. It looks like those Bs were looking for an excuse to switch from what *ought to be done* to what *is done*.

The same does not apply when anticipations about dishonesty are disproven. Considering the last two bars in Figure 1, Bs who anticipated A to lie and whose expectations were confirmed, lied on average 2.23 (fourth bar); Bs who anticipated A to lie and whose expectations were not confirmed lied on average 1.5 (third bar), but the difference is not significant. This suggests that honesty does not pass through as an example to those who anticipated their pairs to lie. At least not to the degree that dishonesty does, as it is important to note that, despite the insignificance of the difference, the direction of the effect is consistent with some transmission. It seems that if one is set to lie, observing another who did not, does not make social image concerns, or shame and guilt more salient, or, even if it does, these motives are not strong enough here to make individuals switch to the path of honesty.

Although it is clear that Hypothesis 2A is confirmed but 2B is not, we engage in further analysis to evaluate the level of transmission of dishonesty but without considering Bs beliefs about As' behavior. In Figures 2a and 2b below we depict the proportion of Bs who lied after observing an honest or dishonest A and the average lying of Bs in the equivalent cases.

Figure 2a: Proportion of Dishonest Bs over Honest & Dishonest As

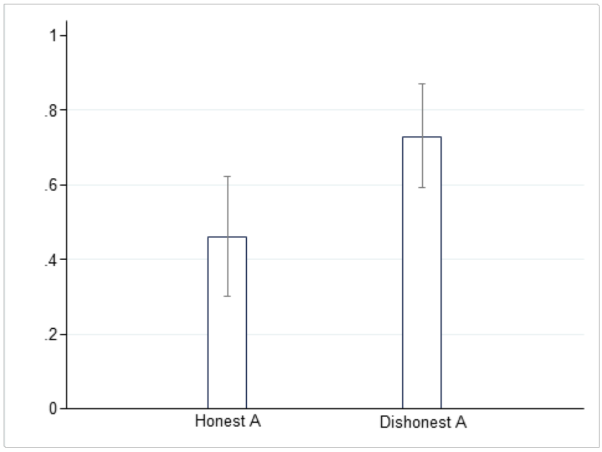

Figure 2b: Average lying of Bs over Honest & Dishonest As

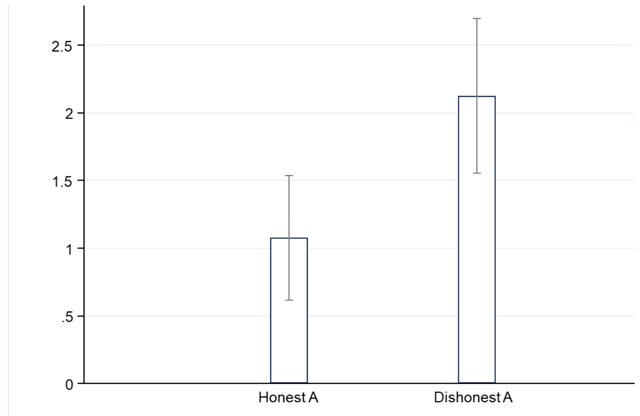

In Figure 2a the bars stand for fractions of subjects, in Figure 2b the bars stand for average lying, and the spikes represent 95% confidence intervals. In Figure 2a, we see that the fraction (46,15%) of dishonest Bs who observe an honest A is significantly lower from the fraction (72.5%) of dishonest Bs who observe a dishonest A with $p=0.01$. In Figure 2b, the average lying of Bs who observe an honest A (1.07) is significantly lower from the average lying of those Bs who observe a dishonest A (2.12) with $p < 0.01$. It is clear that lying occurs more frequently and at higher levels when individuals observe others lying (regardless of their beliefs).

5. Conclusion

In this paper, we presented evidence on the transmission of honesty and dishonesty using a laboratory experiment. In particular, we evaluated whether individuals faced with a lying dilemma are willing to switch from their intended behavior to the behavior of a peer they observed, when intentions and observations do not overlap. The results point to important asymmetries. While individuals who were likely to act honestly switch to acting dishonestly when observing a dishonest peer, individuals who were likely to act dishonestly do not switch to acting honestly when observing an honest peer. In other words, dishonesty seems to transmit while honesty does not. It is a frustrating result. Why is antisocial but not prosocial behavior easy to catch on? If anything, this is not in line with social conformity and imitation theories. These theories would predict that followers follow independently of the moral connotation of the action being followed. Similarly, theories of social image put forward by the lying and deception literature would predict that image concerns, shame, and guild would work as lubricants for honesty to transmit when one observes another, who in turn will witness own actions, acting honestly.

We are not the first to report such asymmetries. Our results are in line with Lefebvre et al., 2014 who study tax compliance behavior. They report that examples of high compliance do not produce a disciplining effect, but examples of low compliance significantly increase tax evasion. Such findings might be more common in non-social settings. The present study was not designed for, and cannot bring answers to why, socially harmful behaviors transmit while socially beneficial behaviors do not. If any conclusions can be drawn, they would point to the need for further investigation of the matter. More theories, experiments, and empirical exercises need to be performed to verify the replicability and the generalizability of the asymmetries reported here. If these are deemed robust, then further research will be needed to look into the mechanisms that bring about these results. Until then, a reserved recommendation to the interested policy makers would be to work on building contextual choice architectures that limit observability of peers' actions in matters of significant moral hypostasis.

Acknowledgments: We are thankful to Adrian Bruhin, Marina Chugunova, Rustamdjan Hakimov, Stephanie Heger, Friederike Mengel, Nikos Nikiforakis, Ernesto Reuben, Simeon Schudy, Joel Sobel, Lisa Spantig, Marie Claire Villeval, and Christian Zehnder.

Declarations:

Funding The study was financed by the personal research budgets of author 3 from the University of XX, and of author 2 from the XX University XX.

Competing Interests The authors declare no competing interests

Availability of data Data will become publicly available upon publication, and will be available upon request until then.

Ethics approval The study received ethics approval from the University of XX and IRB from XX University (#0742018).

References

Abeler, Johannes, Anke Becker, and Armin Falk. 2014. "Representative evidence on lying costs." *Journal of Public Economics* 113: 96-104. doi:dx.doi.org/10.1016/j.jpubeco.2014.01.005.

Abeler, Johannes, Daniele Nosenzo, and Collin Raymond. 2019. "Preferences for truth-telling." *Econometrica* 1115-1153.

Barr, Abigail , and Georgia Michailidou. 2017. "Complicity without Connection." *Journal of Economic Behavior & Organization* 142: 1-10. doi:https://doi.org/10.1016/j.jebo.2017.07.013.

Cialdini, Robert B., Raymond R. Reno, and Carl A. Kallgren. 1990. "A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places." *Journal of Personality and Social Psychology* 1015-1026.

Conrads, J. , B. Irlenbusch, R. M. Rilke, and G. Walkowitz. 2013. "Lying and team incentives." *Journal of Economic Psychology* 1-7.

Croson, Rachel, and Yen Shang. 2008. "The impact of downward social information on contribution decisions." *Experimental Economics.* doi:DOI 10.1007/s10683-007-9191-z.

Dana, Jason, Roberto A. Weber, Jason Xi Kuang, and . 2007. "Exploiting moral wiggle room: experiments." *Economic Theory* 33: 67-80. doi:10.1007/s00199-006-0153-z.

Diekmann, Andreas, Wojtek Przepiorka, and Heiko Rauhut. 2015. "Lifting the veil of ignorance: An experiment on the contagiousness of norm violations." *Rationality and Society* 309-333.

Ederer, Florian, and Ernst Fehr. 2007. "Deception and incentives: How dishonesty undermines effort provision." *University of Zurich Institute for Empirical Research in Economics Working Paper* 341.

Erat, Sanjiv, and Uri Gneezy. 2012. "White lies." *Management Science* 58 (4): 723-733.

Fischbacher, Urs, and Franziska Föllmi-Heusi. 2013. "Lies in disguise—an experimental study on cheating." *Journal of the European Economic Association* 11 (3): 525-547.

Fortin, B., Lacroix, G., and M. C. Villeval. 2007. "Tax evasion and social interactions." *Journal of Public Economics* 2089-2112.

Gächter, Simon, and Elke Renner. 2018. "Leaders as role models and 'belief managers' in social dilemmas." *Journal of Economic Behavior and Organization.* https://doi.org/10.1016/j.jebo.2018.08.001.

Gächter, Simon, and Jonathan F Schulz. 2016. "Intrinsic honesty and the prevalence of rule violations across societies." *Nature* 531: 496-499.

Gino, Francesca, Shahar Ayal, and Dan Ariely. 2009. "Contagion and Differentiation in Unethical Behavior: The Effect of One Bad Apple on the Barrel." *Psychological Science* 393-398.

Gneezy, Uri. 2005. "Deception: The role of consequences." *The American Economic Review* 95 (1): 384-394.

Gneezy, Uri, Agne Kajackaite, and Joel Sobel. 2018. "Lying avesion and the size of the lie." *American Economic Review* 419-53.

Gneezy, Uri, Bettina Rockenbach, and Marta Serra-Garcia. 2013. "Measuring Lying Aversion." *Journal of Economic Behavior & Organization.*

Hugh-Jones, David, and Jinnie Ooi. 2017. "Where do fairness preferences come from? Norm transmission in a teen friendship network." *University of East Anglia School of Economics Working Paper Series 2017-02.* https://ueaeco.github.io/working-papers/papers/ueaeco/UEA-ECO-17-02.pdf.

Hurkens, Sjaak,, and Navin Kartik. 2009. "Would I lie to you? On social preferences and lying aversion." *Experimental Economics* 12 (2): 180-192. doi:10.1007/s10683-008-9208-2.

Kajackaite, Agne, and Uri Gneezy. 2017. "Incentives and cheating." *Games and economic behavior* 433-444.

Kartik, N, O Tercieux, and R Holden. 2014. "Simple mechanisms and preferences for honesty." *Games and Economic Behavior* 284-290.

Keizer, Kees, Siegwart Lindenberg, and Linda Steg. 2008. "The spreading of disorder." *Science* 1681-1685.

Kocher, Martin G., Simeon Schudy, and Lisa Spantig. 2016. "I Lie? We Lie! Why? Experimental Evidence on Dishonesty Shift in Groups." *Management Science* 3995-4008.

Lauer, T., and A Untertrifaller. 2019. "Conditional Dishonesty." *mimeo.*

Lefebvre, Mathieu, Pierre Pestieau, Arno Riedl, and Marie Claire Villeval. 2015. "Tax evasion and social information: an experiment in Belgium, France, and the Netherlands." *International tax and public finance* 401-425.

Levine , Emma E., and Maurice E. Schweitzer. 2014. "Are liars ethical? On the tension between benevolence and honesty." *Journal of Experimental Social Psychology* 107–117.

López-Pérez, R, and E Spiegelman. 2013. "Do Economists Lie More?" *Experimental Economics* 233-247.

Mazar, Nina, On Amir, and Dan Ariely . 2008. "More ways to cheat-expanding the scope of dishonesty." *Journal of Marketing Research* 45 (6): 651-653.

Rilke, Rainer Michael, Anastasia Danilov, Ori Weisel, Shaul Shalvi, and Bernd Irlenbusch. 2021. "When leading by example leads to less corrupt collaboration." *Journal of Economic Behavior and Organization.* doi:https://doi.org/10.1016/j.jebo.2021.05.007.

Shalvi, S., M. J. J. Handgraaf, and C. K. W. De Dreu. 2011. "Ethical manoeuvering: Why people avoid both major and minor lies." *British Journal of Management.*

Shalvi, Shaul, Ori Eldar, and Yoella Bereby-Meyer. 2012. "Honesty Requires Time (and Lack of Justifications)." *Psychological Science.*

Shang, Jen, and Rachel Croson. 2009. "A Field Experiment in Charitable Contribution: The Impact of Social Information on the Voluntary Provision of Public Goods." *The Economic Journal.* https://doi.org/10.1111/j.1468-0297.2009.02267.x.

Smerdon, David, Theo Offerman, and Uri Gneezy. 2020. "'Everybody's doing it': on the persistence of bad social norms." *Experimental Economics* 392- 440. doi:https://doi.org/10.1007/s10683-019-09616-z.

Thöni, Christian, and Simon Gächter. 2015. "Peer effects and social preferences in voluntary cooperation: A theoretical and experimental analysis." *Journal of Economic Psychology.* https://doi.org/10.1016/j.joep.2015.03.001.

Weisel, Ori, and Shaul Shalvi. 2015. "On the collaborative roots of corruption." *PNAS* 10651-10656.

Appendix

Section A.1. Experimental Instructions
(Translated from French)

General Instructions ***(common for simultaneous & sequential)***

You are about to participate in an economic experiment. The experiment is conducted by the Department of Economics of the University of Lausanne.

For your participation in the experiment you will earn a payment of CHF 10 for sure. The experiment allows you to earn additional money. At the end of the experiment, you will be paid CHF 10 and any additional money you earned during the experiment. It is to your own benefit to read these explanations carefully.

You can perform the experiment at your own speed. It is prohibited to communicate with the other participants during the whole course of the experiment. If you do not abide by this rule you will be excluded from the experiment and all payments. However, if you have questions you can always ask one of the experimenters by raising your hand.

You can abort the experiment anytime you wish without giving any reasons. To do so, please raise your hand and tell the experimenters that you wish to abort the experiment. One experimenter will then guide you outside the laboratory. You are not eligible to any payments in case you abort the experiment.

Your anonymity is guaranteed.

At the end of the experiment, one experimenter will give you a payment sheet with the amount you will be paid. You will need to carry the payment sheet with you and present it to an experimenter outside the LABEX. The experimenter outside the LABEX does not know about any of the decisions you made during the experiment. This experimenter will then pay you according to your payment sheet. After that you will sign a form stating that you received the payment. Since the form you sign does not contain your participant number, there is no way any experimenter can determine your identity.

If you have any questions right now, please raise your hand. Otherwise, you can now proceed with the detailed explanations of the experiment.

Thank you very much for your participation!

Specific Instructions –*Simultaneous-*

We are now going to explain the task you will perform. For this task you are randomly and anonymously paired with another participant in this room. One participant is randomly assigned to the role of Person A and the other participant to the role of Person B. You will learn whether you have been assigned to the role of Person A or Person B in the end of these instructions. Person A and Person B will observe the outcome of an electronic six sided die-roll. Both Person A and Person B will observe the same outcome of this die-roll. The experimenter will also observe the outcome of this die-roll. This die-roll has six possible outcomes: 1, 2, 3, 4, 5, and 6.

Each outcome can be realized with a probability of 1/6. The table below summarizes the die-roll outcomes and their associated probabilities.

Outcome of the die-roll	1	2	3	4	5	6
Probability	1/6	1/6	1/6	1/6	1/6	1/6

Your task, and the other person's task, is the same: to report the outcome of the die-roll. Each person's monetary payment is only determined by his/her report. In other words, Person A's monetary payment is only determined by Person A's report and Person B's monetary payment is only determined by Person B's report. If you report number 1, then you are paid CHF 1, if you report number 2, then you are paid CHF 2, etc. Here is a table of how your report is associated with your monetary payment:

Your report	1	2	3	4	5	6
Monetary Payment in CHF	1	2	3	4	5	6

Once Person A and Person B have made their reports, each will be asked to make a guess. This will become clear during the experiment. If Person A's guess is correct, Person A will be paid an additional CHF 1. If Person B's guess is correct, Person B will be paid an additional CHF 1. Finally, Person A observes the report of Person B and Person B observes the report of Person A.

Therefore, the sequence of this experiment is as follows:

1. Person A and Person B observe the outcome of the die-roll
2. Person A and Person B make their reports
3. Person A and Person B make their guesses
4. Person A and Person B observe each other's reports
5. Person A is paid his/her report and Person B is paid his/her report

The three examples that follow should make it clear how Person A's report and Person B's report are related to the monetary payments in this experiment.

Example 1: Assume the outcome of the die-roll is 4, Person A reports 5, and Person B reports 4. In this example, Person A is paid CHF 5 and Person B is paid CHF 4.

Example 2: Assume the outcome of the die-roll is 2, Person A reports 4, and Person B reports 5. In this example, Person A is paid CHF 4 and Person B is paid CHF 5.

Example 3: Assume the outcome of the die-roll is 3, Person A reports 3, and Person B reports 3. In this example, Person A is paid CHF 3 and Person B is paid CHF 3.

It is important that you have a good understanding of the experimental instructions. To check that the instructions are clear to you we now ask you to answer a few questions. Your answers to these questions do not have any influence on the experiment itself or on the payment you will receive at the end of the experiment. The experiment will start once you and the person you are paired with have answered the questions correctly.

Questions to check your understanding:

1. If you report 5, how much are you paid?
2. If the outcome of the die-roll is 3 and you report 2, how much are you paid?
3. Does Person A observe a different outcome of the die-roll than Person B? Yes or No?
4. If the outcome of the die-roll is 2 and the person you are paired with reports 3, how much is the person you are paired with paid?
5. Does the report of one person influence the monetary payment of the other person? Yes or No?

Specific instructions –Sequential-

We are now going to explain the task you will perform. For this task you are randomly and anonymously paired with another participant in this room. One participant is randomly assigned to the role of Person A and the other participant to the role of Person B. You will learn whether you have been assigned to the role of Person A or Person B in the end of these instructions.

Person A and Person B will observe the outcome of an electronic six-sided die-roll. Both Person A and Person B will observe the same outcome of this die-roll. The experimenter will also observe the outcome of this die-roll. This die-roll has six possible outcomes: 1, 2, 3, 4, 5, and 6. Each outcome can be realized with a probability of 1/6. The table below summarizes the die-roll outcomes and their associated probabilities.

Outcome of the die-roll	1	2	3	4	5	6
Probability	1/6	1/6	1/6	1/6	1/6	1/6

Your task, and the other person's task, is the same: to report the outcome of the die-roll. Each person's monetary payment is only determined by his/her report. In other words, Person A's monetary payment is only determined by Person A's report and Person B's monetary payment is only determined by Person B's report. If you report number 1, then you are paid CHF 1, if you report number 2, then you are paid CHF 2, etc. Here is a table of how your report is associated with your monetary payment:

Your report	1	2	3	4	5	6
Monetary Payment in CHF	1	2	3	4	5	6

Once Person A and Person B have observed the outcome of the die-roll, Person A makes his/her report. After this, Person B observes Person A's report. Then, Person B makes his/her report. Finally, Person A observes Person B's report.

During the experiment Person A and Person B will also be asked to make a guess. This will become clear during the experiment. If Person A's guess is correct, Person A will be paid an additional CHF 1. If Person B's guess is correct, Person B will be paid an additional CHF 1.

Therefore, the sequence of this experiment is as follows:

1. Person A and Person B observe the outcome of the die-roll
2. Person A makes his/her report
3. Person A and Person B make their guesses
4. Person B observes Person A's report
5. Person B makes his/her report
6. Person A observes Person B's report
7. Person A is paid his/her report and Person B is paid his/her report

The three examples that follow should make it clear how Person A's report and Person B's report are related to the monetary payments in this experiment.

Example 1: Assume the outcome of the die-roll is 4, Person A reports 5, and Person B, after having observed Person A's report, reports 4. In this example, Person A is paid CHF 5 and Person B is paid CHF 4.

Example 2: Assume the outcome of the die-roll is 2, Person A reports 4, and Person B, after having observed Person A's report, reports 5. In this example, Person A is paid CHF 4 and Person B is paid CHF 5.

Example 3: Assume the outcome of the die-roll is 3, Person A reports 3, and Person B, after having observed Person A's report, reports 3. In this example, Person A is paid CHF 3 and Person B is paid CHF 3.

It is important that you have a good understanding of the experimental instructions. To check that the instructions are clear to you we now ask you to answer a few questions. Your answers to these questions do not have any influence on the experiment itself or on the payment you will receive at the end of the experiment. The experiment will start once you and the person you are paired with have answered the questions correctly

Questions to check your understanding:

1. If you report 5, how much are you paid?
2. If the outcome of the die-roll is 3 and you report 2, how much are you paid?
3. Does Person A observe a different outcome of the die-roll than Person B? Yes or No?
4. Who is the first person to report the outcome of the die-roll? Person A or Person B?
5. Does the report of one person influence the monetary payment of the other person? Yes or No?

Appendix A.2. Ztree screenshots and demographic questions

Below we provide a few snips of the screens as seen by participants and we provide the English translation below each.

Screenshot 1: You have been assigned the role of Person B

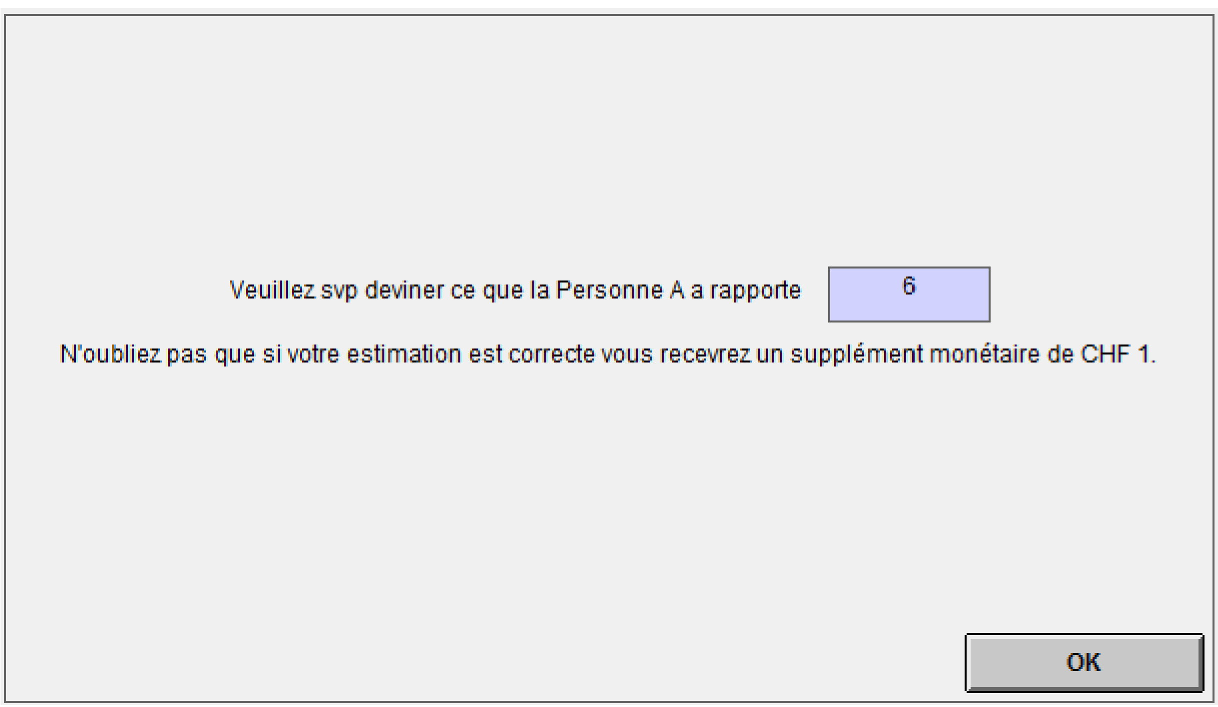

Screenshot 2: Please guess the report of Person A. Recall that if you guess right you will receive 1 CHF in addition

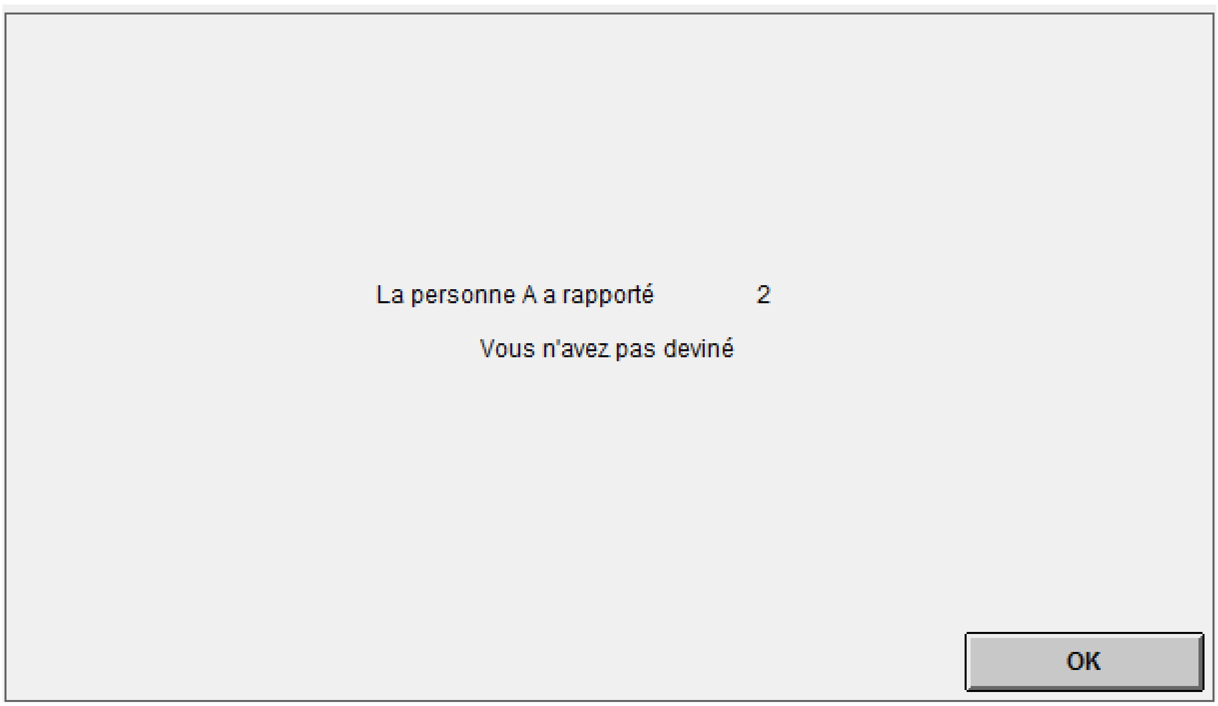

Screenshot 3: Person A reported 2. You have not guessed

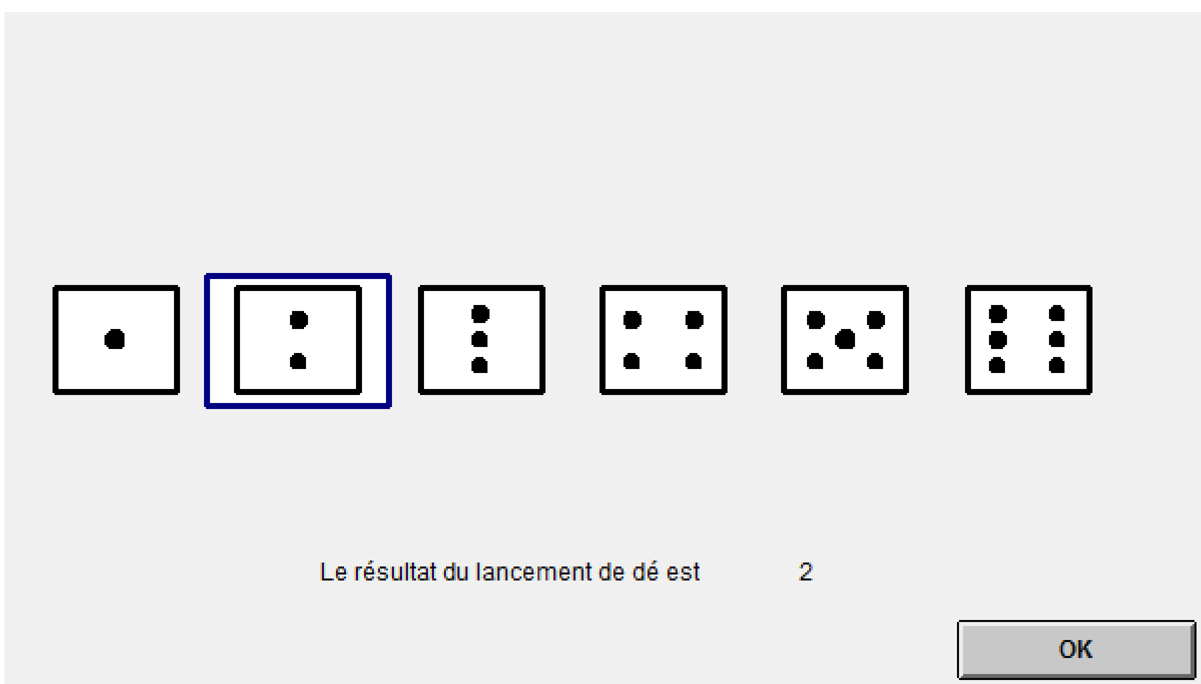

Screenshot 4: The outcome of the die roll is 2

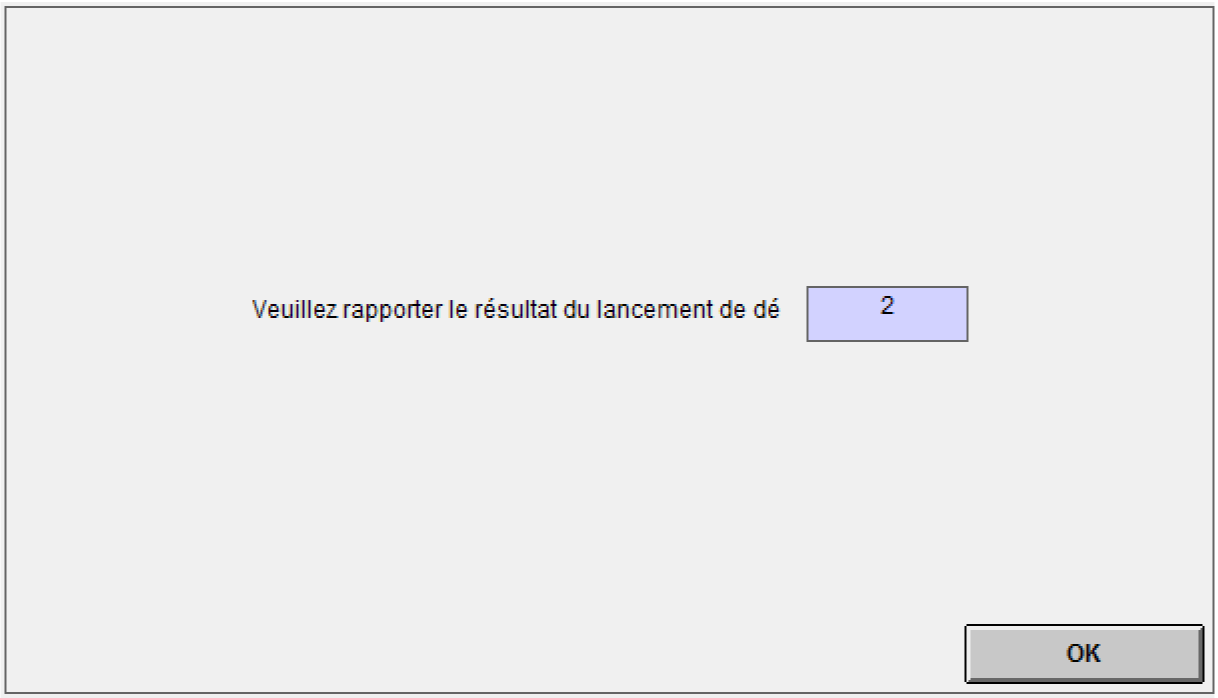

Screenshot 5: Please report the outcome of the die roll (here, a subject reports 2)

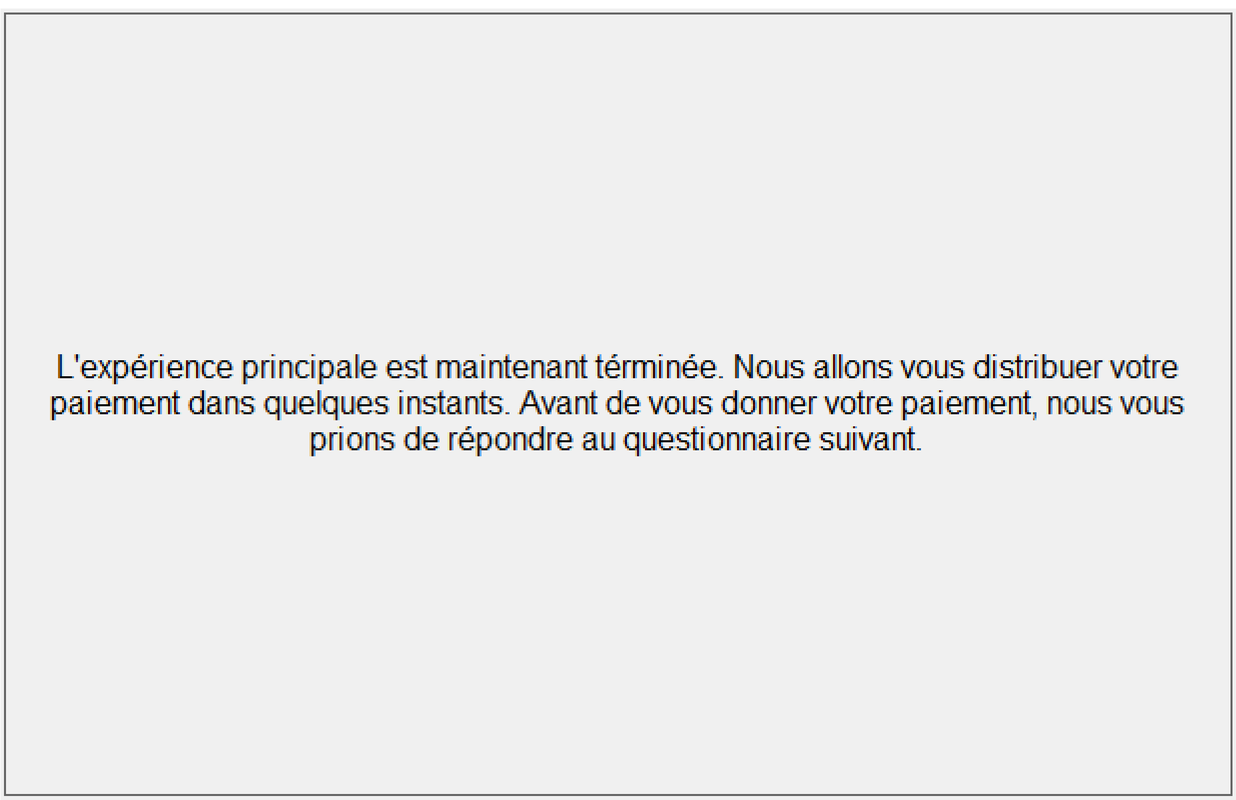

Screenshot 6: The experience is now over. We will soon give you your payment. Before giving you your payment, we kindly ask you to fill in the following questionnaire.

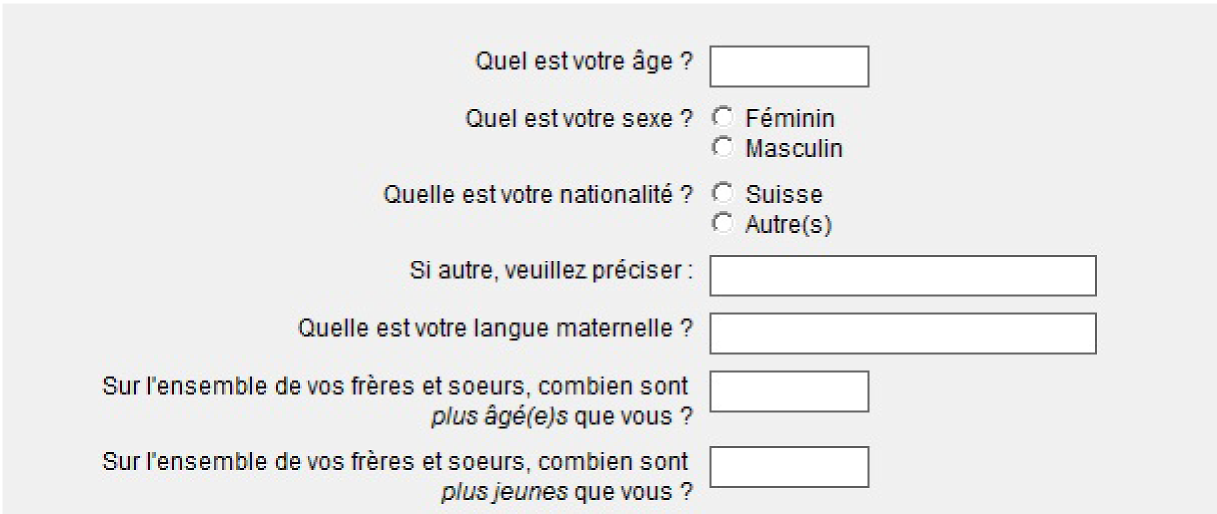

Quel est votre âge ?

Quel est votre sexe ? Féminin / Masculin

Quelle est votre nationalité ? Suisse / Autre(s)

Si autre, veuillez préciser :

Quelle est votre langue maternelle ?

Sur l'ensemble de vos frères et soeurs, combien sont *plus âgé(e)s* que vous ?

Sur l'ensemble de vos frères et soeurs, combien sont *plus jeunes* que vous ?

Screenshot 7: How old are you? What is your gender? Which is your nationality? Swiss/other - If other, please specify - What is your mother tongue? – Among your siblings, how many are older than you are? - Among your siblings, how many are younger than you?"

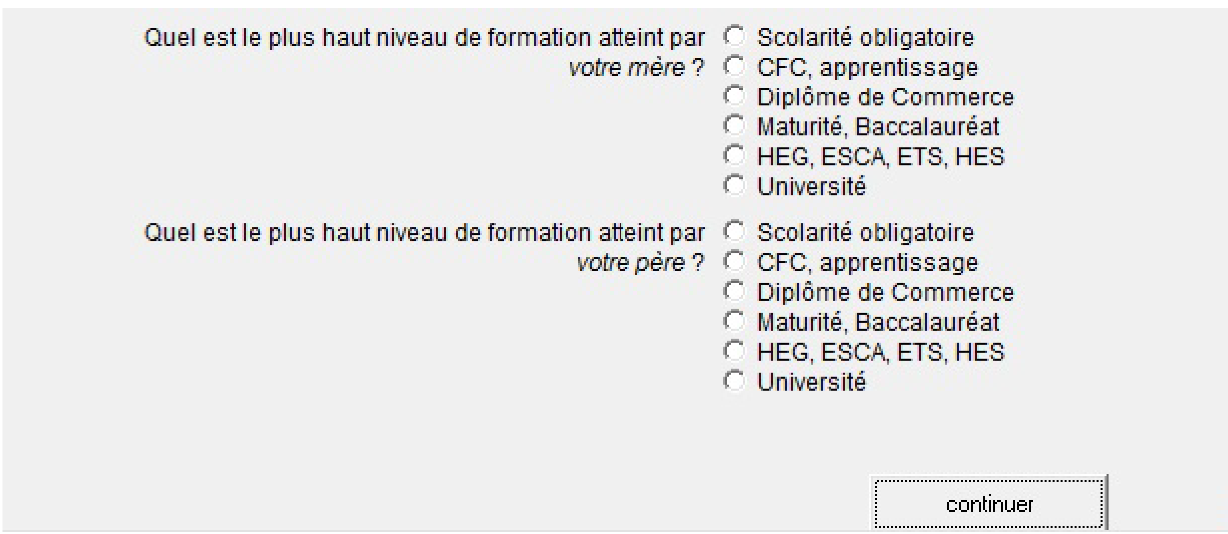

Quel est le plus haut niveau de formation atteint par *votre mère* ?
- Scolarité obligatoire
- CFC, apprentissage
- Diplôme de Commerce
- Maturité, Baccalauréat
- HEG, ESCA, ETS, HES
- Université

Quel est le plus haut niveau de formation atteint par *votre père* ?
- Scolarité obligatoire
- CFC, apprentissage
- Diplôme de Commerce
- Maturité, Baccalauréat
- HEG, ESCA, ETS, HES
- Université

continuer

Screenshot 8: Which is the higher education level obtained by your mother? Compulsory school/ Internship / Vocational Education / High school / Bachelor degree / Master degree - Which is the higher education level obtained by your father? Compulsory school/ Internship / Vocational Education / High school / Bachelor degree / Master degree.

Combien de participants à cette séance connaissiez-vous déjà avant l'expérience ?

Combien d'habitants compte la commune dans laquelle vous avez habité pour la plus grande partie de votre vie ?

- jusqu´à 2 000 habitants
- 2 000 à 10 000 habitants
- 10 000 à 100 000 habitants
- plus de 100 000 habitants

Screenshot 9: Among the participants to this experience, how many people did you already know? - How many inhabitants are there in the city you have spent most of your life in? up to 2000 / among 2000 and 10000 / among 10000 and 100000 / more than 100000

Où faites-vous vos études ?

- EHL - Ecole Hôtelière de Lausar
- EPFL - Environnement naturel,
- architectural et construit
- EPFL - Informatique
- EPFL - Sciences de la vie
- EPFL - Sciences et techniques c
- EPFL - Sciences de base
- EPFL - Collège des humanités
- EPFL - Collège du management
- HEP - Haute École Pédagogique
- HES - Haute École de la Santé
- UNIL - Faculté de Biologie et de
- UNIL - Faculté de Droit et des Sc
- UNIL - Faculté de Théologie
- UNIL - Faculté des Géosciences
- UNIL - Faculté des Hautes Etude
- UNIL - Faculté des Lettres
- UNIL - Faculté des Sciences So
- Autre
- Je ne suis pas étudiant-e

Screenshot 10: Where are you currently studying?

Quelle est votre année d'études ?

- 1ère Bachelor
- 2ème Bachelor
- 3ème Bachelor
- 1ère Master
- 2ème Master
- Autre
- Je ne suis pas étudiant-e

Quelle est la moyenne (des notes) que vous avez obtenu lors de votre dernière année (soit à l'université, soit au gymnase, soit une autre institution éducative) ? *Si vos notes proviennent d'une instutition en dehors de la Suisse, pourriez-vous, si possible, convertir votre moyenne obtenu à l'étranger au système de note Suisse allant de 1 à 6.*

OK

Screenshot 11: At which year of your current studies are you?

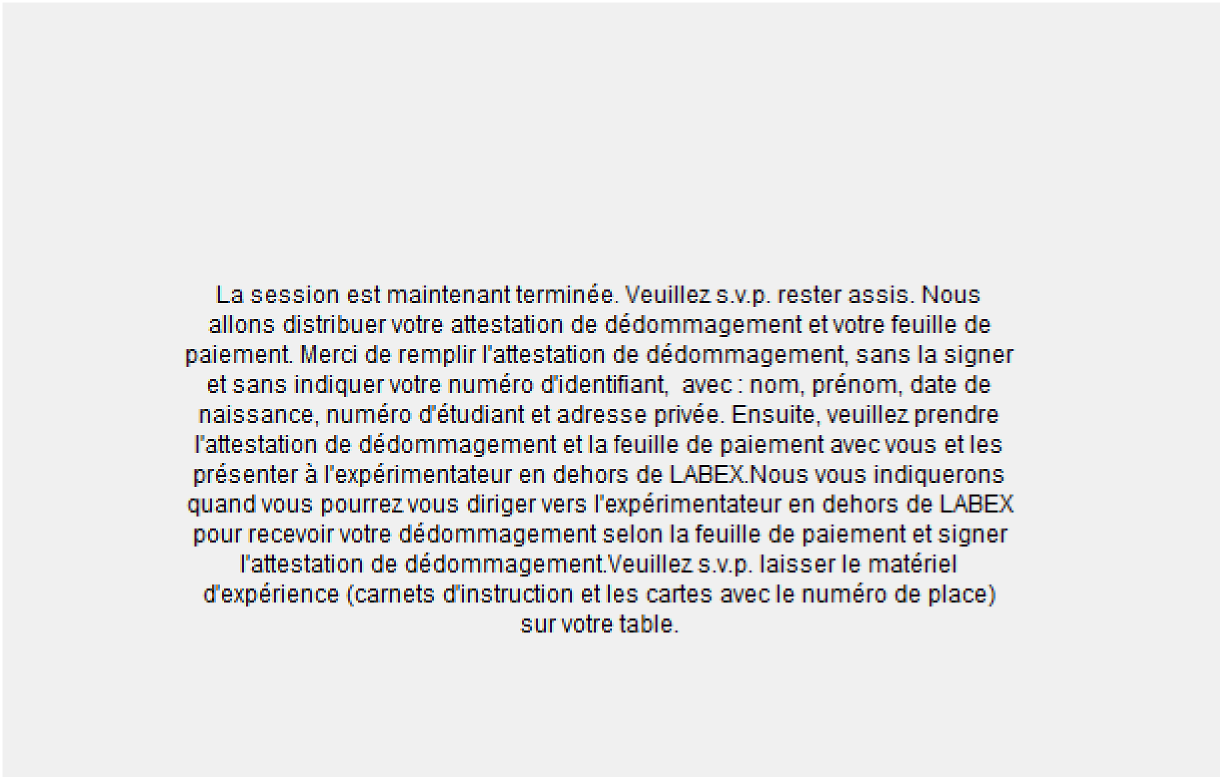

Screenshot 12: The experience is now over. Please remain seated. We are about to distribute you your payment sheet. Please fill the payment receipt, without signing it nor indicating your subject number, with your name, surname, student ID and address. Then please take the payment sheet and the payment receipt with you and show them to the experimenter outside of LABEX. We will tell you when to leave the room and reach the experimenter outside of LABEX to receive your compensation according to the payment sheet and sign the payment receipt. Please leave the experiment material (instructions and cubicle number cards) on your desk.

Appendix B. The Control Treatment and its Results Compared to Gneezy et al.,2018.

Control Treatment Design

Similarly to the sequential and simultaneous treatments, in the control treatment subjects enter the lab, they are paired, and are assigned the role of A or B in a random fashion. Here, however, Bs are passive and they will not be engaging in reporting. Although both A and B in each pair observe a die roll outcome, only A will be reporting and his/her report will determine his/her payoff. Once A reports, B (whose presence is meant to trigger similar audience effects as in the treatments) will observe A’s report. B’s payoff is determined as follows. Before subjects get paired and learn about their tasks, we describe to them the simultaneous treatment (in non-indicative language) which we call "Experiment 1". We explained that some students, from the same university as themselves, had already participated in "Experiment 1" and that we have kept their reports. Then, we explained that B's payoff would be determined by picking one of the reports of the subjects in "Experiment 1" who had observed the same die-roll outcome. For example, if a pair in the control observed a 3, A would receive a payoff equal to his/her report, and B would receive a payoff equal to the report of a B who participated in "Experiment 1" and had observed a die-roll of 3.

We decided to go with this rather non-straightforward control treatment for the following reasons. First, if in the sequential and the simultaneous treatments subjects' reports are observed by each other, then the same audience effects must apply for subjects in the control treatment. Second, passive Bs' payoffs in the control would have to be such that the induced other regarding preferences of As are constant across the control and the treatments. If, for example, we paid passive Bs in the control with a flat fee, inequality averse As might have been anchored to this flat fee and report according to it. In total, 138 subjects participated in the control but only half of them were active (As), the mean age of those half was 21 years, and 52.17% were female.

Comparison of Control Results with Gneezy et al., 2018

We will qualitatively compare the results of our control to the lying levels reported in Gneezy et al., 2018 (GKS hereafter). In particular, we will be using the observed number treatment of that paper for a qualitative comparison. We will be using this paper as a reference for various reasons. First, although GKS do not use a die rolling game as we do, our designs are mostly comparable. Second, lying in GKS is observable by the experimenter as is in our case; this is important as we can argue that the audience effects experienced by the subjects of both experiments are similar. Third, in both experiments, subjects lie for own profit, without strategic interaction and without payoff complementarities between groups of subjects. Fourth, both experiments were executed in similar countries (Germany and Switzerland) therefore, any cultural component that might be confounding with morality should not be an issue in this comparison. Table below provides a summary of our results and the GKS results.

	This Experiment (Control)	GKS

% of Liars across all subjects	44.93%	26.41%
% of Maximal liars among those that lied	64.52%	67.96%
Average lying (normalized)	0.19	0.13

To an extent, our results are replicating GKS. Although in the extensive margin, there are more liars in our control, there seems to be a consistency of behavior in the intensive margin. Among those that lied, the percentage of those who lied maximally is comparable in the two experiments. As the lying opportunity scale varies across the two experiments (5 units in this experiment, 9 in the GKS) we proceeded with normalization and we observe that average lying is also comparable across the two experiments. We normalized by clustering all reports according to the associated observation and we divided each report with the available lying interval given the observation. For example, considering subject(i) who observed (b) and reported (r), we normalized by using (r-b)/(6-b) for our experiment, and (r-b)/(10-b) for the GKS experiment, and then we proceeded by averaging the normalized lying. Given these results, we can conclude that we are very close in replicating GKS results.

Appendix C. A Graphical Summary of Results

We use the section below to present a different graphical summary of our results. We start with the distribution of reported payoffs across the three experimental conditions and we proceed with a graph of percentage of liars conditional on the actual payoff. Then, we present the line of the percentage of liars conditional on different observations. Last, we provide a breakdown of maximal and partial liars among our three experimental conditions.

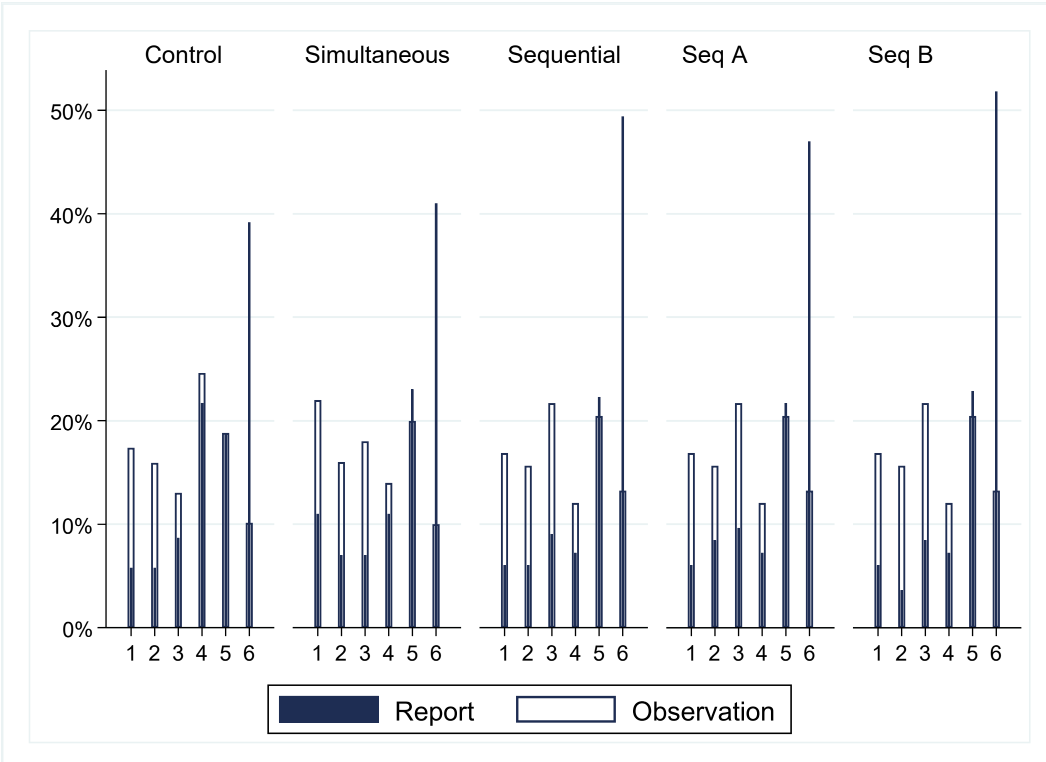

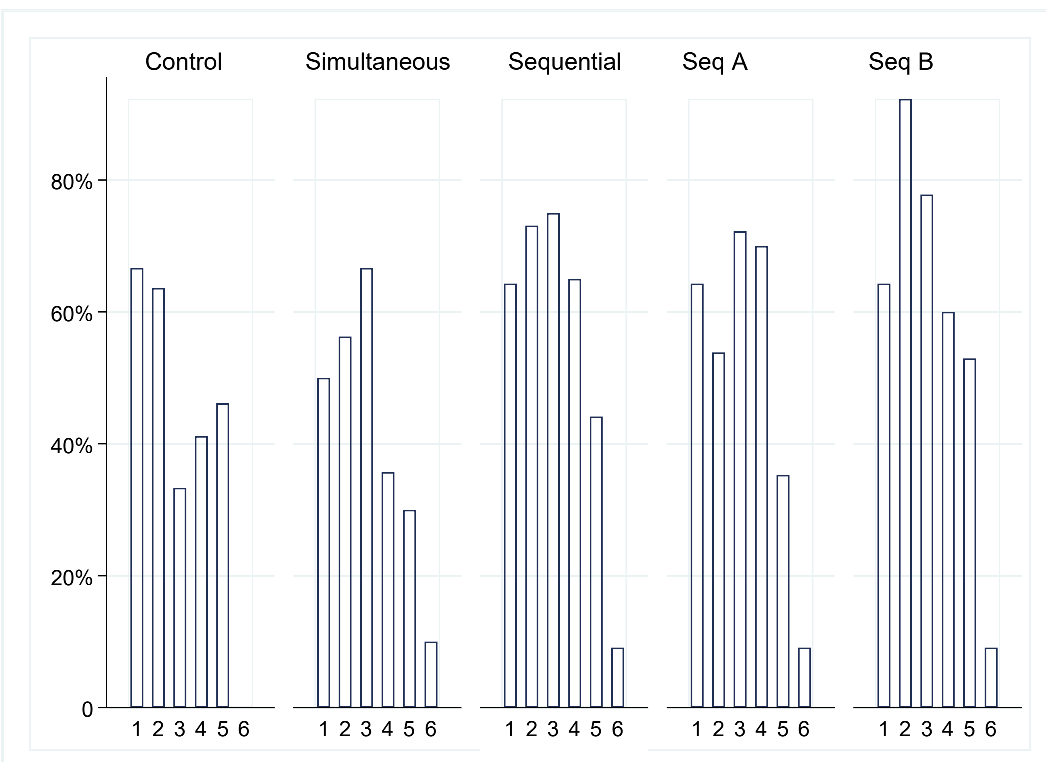

Appendix Figure 2 Fraction of Liars Conditional on the Actual Payoff

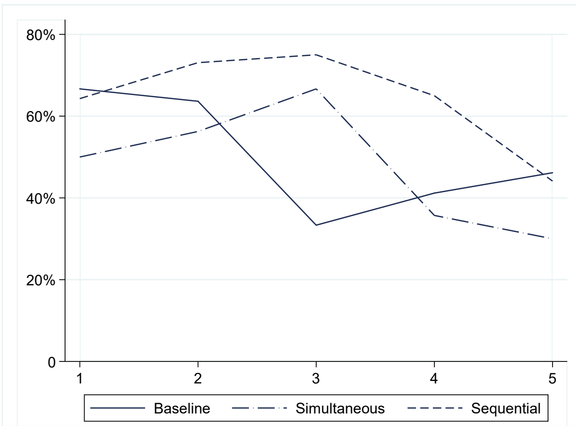

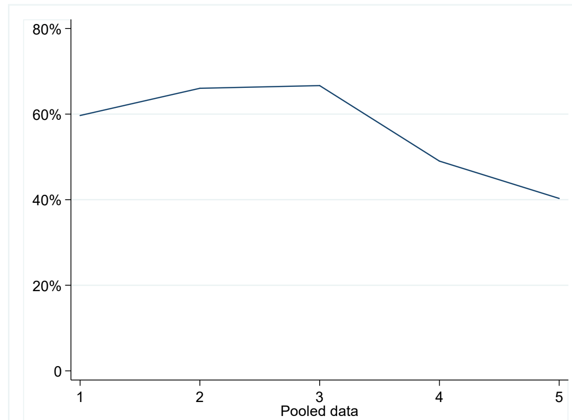

Appendix Figure 3 Line of % of liars conditional on observation

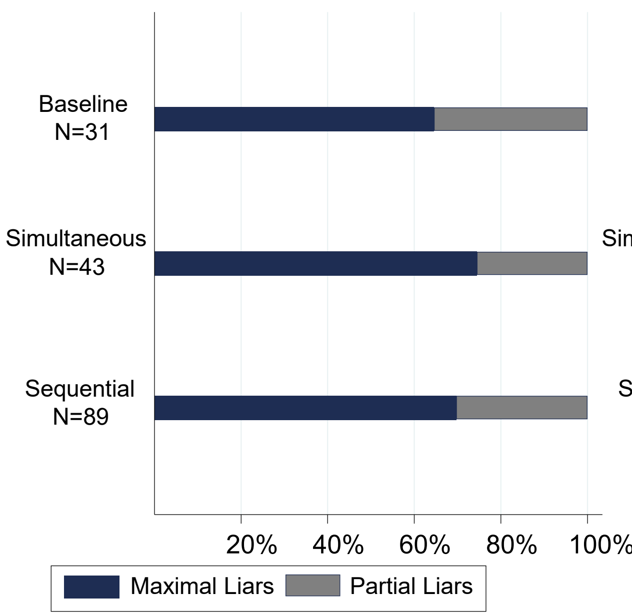

Figure 4: Maximal & Partial Lying among Liars

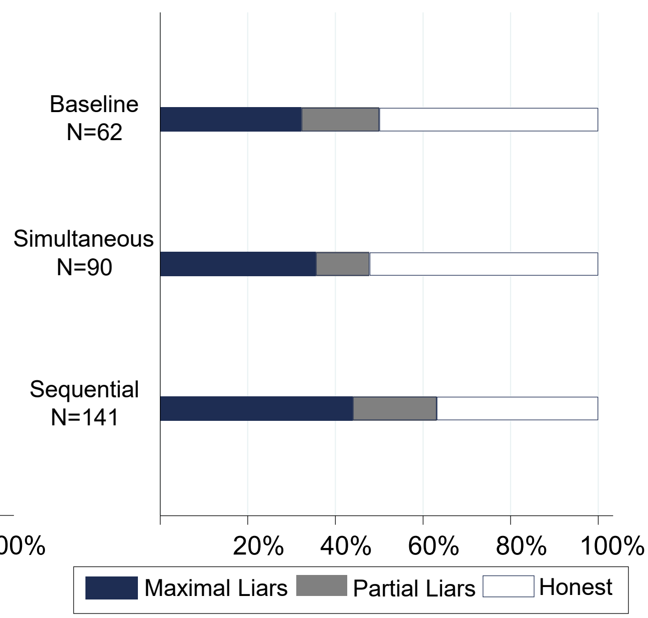

Figure 5: Maximal - Partial Lying & Honesty among Subjects

Appendix D. Regressions

	(1) LYING OF B	**(2)** LYING OF B	**(3)** LIAR B DUMMY	**(4)** LIAR B DUMMY	**(5)** LYING OF B	**(6)** LYING OF B	**(7)** LYING OF B	**(8)** LYING OF B	**(9)** LIAR B DUMMY	**(10)** LYING OF B	**(11)** LYING OF B
LIAR A DUMMY	1.04 (0.006)	0.67 (0.09)	0.26 (0.01)	0.10 (0.37)					0.04 (0.67)		
LYING OF A					0.44 (0.001)	0.37 (0.007)	0.188 (0.07)	0.18 (0.10)		0.08 (0.48)	
B'S GUESS IF A IS LIAR DUMMY									0.40 (0.01)		
BS' GUESS OF AS' LYING										0.45 (0.004)	0.54 (0.001)
DIF. BS' GUESS & AS' LYING											-0.08 (0.48)
INCLUDES THOSE OBSERVING 6	Yes	No	Yes	No	Yes	No	Yes	No	No	No	No
N	79	69	79	69	79	69	79	69	69	69	69
FIXED EFFECTS	No	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes
CONSTANT	1.07 (<0.001)	1.44 (<0.001)	0.46 (<0.001)	0.62 (<0.001)	0.99 (<0.001)	1.26 (<0.001)	0.52 (0.25)	0.52 (0.28)	0.28 (0.04)	0.27 (0.04)	0.27 (0.04)

The regressions bellow present results from the sequential treatment only. In the table, the parenthesis present p-values.

Regression (1) is a linear regression of the lying of B on a liar A dummy which equals 1 when A lies and 0 otherwise. We see that when A does not lie, B lies by 1.07 (given by the constant), while when A lies, there is an increase of 1.04 in B's lying. This coefficient is significant at a 1% level. To see the probability of B lying after observing a dishonest A, we regress a liar B dummy, on the liar A dummy as shown in linear regression (3). We see that the probability of B lying if A is honest is approximately 46% (given in the constant) and it increases by approximately 26% if A is dishonest which is significant in a 5% level. This means that a B who observes a dishonest A, will be dishonest approximately 7 out of 10 times. In linear regression (5), we see that Bs lie (given by the constant) is 0.99 which means that Bs lie approximately that much when As do not lie.

The coefficient of As' lying, 0.44, is significantly different from zero in a 1% level. The coefficient could be interpreted as the marginal increase of Bs lying in As lying. In other words, Bs' lying increases by 0.44 for every unitary increase in As' lying. For the same regression, when fixed effects are considered (regression 7), The coefficient of As lying, 0.18, still significant at a 10% level, can be interpreted as the increase in the amount of lying of Bs, for every unitary increase in As' lying, for each observation. In regression (9), we see that observing a liar A does not increase the probability that B becomes a liar. However, guessing that an A would be a liar, B has a significant 40% probability of becoming a liar. In regression (10) we see that the size of As lying does not impact the size of lying of Bs. Yet, the guess of B about the size of the lie of A, significantly affects Bs lying. In particular, for every unitary increase in Bs belief about the size of A's lie As, Bs' lying increases by 0.45. In regression (11) the guesses of Bs seem to significantly affect the size of Bs lying, while whether these guesses where accurate does not affect behavior in this context. The last three results suggest that the interplay of guessing and observing might be the most important determinant of behavior in this set up, as we show more analytically in the main body of this manuscript.