



UvA-DARE (Digital Academic Repository)

(Dis)honesty in individual and collaborative settings

A behavioral ethics approach

Leib, M.

Publication date

2021

Document Version

Other version

License

Other

[Link to publication](#)

Citation for published version (APA):

Leib, M. (2021). *(Dis)honesty in individual and collaborative settings: A behavioral ethics approach*.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

**Ethical free-riding:
When honest people find dishonest partners**

This chapter was published as Gross, J*., Leib, M.*, Offerman, T., & Shalvi, S. (2018). Ethical free riding: When honest people find dishonest partners. *Psychological science*, 29(12), 1956-1968. <https://doi.org/10.1177/0956797618796480>

*Shared first authorship.

Humans frequently cooperate to achieve mutual goals. Successful cooperation can lead to better performance, higher innovation, and better overall outcomes (e.g., Rusbult & Van Lange, 2003). Providing people with the freedom to seek trustworthy partners and abandon free-riders helps to develop and safeguard cooperation (Efferson, Roca, Vogt, & Helbing, 2016; Rand, Arbesman, & Christakis, 2011). However, cooperation is also an essential part of corruption (Köbis, van Prooijen, Righetti, & Van Lange, 2016; Weisel & Shalvi, 2015). People often engage in corrupt collaboration—the attainment of personal profits by joint acts of rule violation (Weisel & Shalvi, 2015). Such corrupt collaboration, for example, occurred in the recent Volkswagen scandal, when employees collaboratively manipulated Volkswagen engine software to pass key emissions tests, possibly because of their motivation to obtain team bonuses (Goodman, 2015). These joint rule violations resulted in an estimated societal cost of 59 early deaths and \$450 million (Barrett et al., 2015). Given that partner selection is beneficial for cooperation, could it be that Volkswagen employees' ability to find coworkers willing to manipulate the emissions test results assisted in the emergence and spread of such mutual rule-violating behavior? More generally, can partner selection corrupt?

We experimentally investigated how people's ability to select partners influences the emergence and evolution of corrupt collaboration. Previous studies have revealed that working with other people leads to more dishonesty than working alone (Conrads, Irlenbusch, Rilke, & Walkowitz, 2013; Gino, Ayal, & Ariely, 2013; Sutter, 2009; Weisel & Shalvi, 2015), prosocial lies breed trust (Levine & Schweitzer, 2015), and working with the same partner over time increases bribery as people develop trust (Abbink, 2004). Outside the laboratory, however, people can often choose how long and with whom to interact. Do dishonest people search for a dishonest partner—a “partner in crime”? Do honest people distance themselves from dishonest partners? How does being able to choose a partner influence the contagiousness of dishonesty (i.e., the impact one's dishonesty has on one's partner's dishonesty)? And what is the overall financial impact of corrupt collaboration when people can select their partner versus when they cannot? These are the questions we tackled.

Seeking a partner in crime?

The ability to choose partners allows people who are willing to violate rules (in the present case, lie) to increase personal profit by abandoning those who are not willing to cooperate (Reuben & Stephenson, 2013). Instead, they can search for a like-minded dishonest partner who is willing to violate rules for mutual profit. Hence, we predicted that when given

the opportunity to do so, dishonest people would be more likely to leave their partners when paired with honest versus dishonest people (Hypothesis 1).

But what will honest people who are unwilling to lie for profit do? We consider two possibilities. First, being able to select one's partner allows honest people to avoid engaging in corrupt collaboration. Supporting this possibility, studies have shown that people do not lie much, even when lying is profitable (Abeler, Nosenzo, & Raymond, 2019), they avoid situations that may tempt them to lie (Fishbach, Friedman, & Kruglanski, 2003), and they prefer to interact with those who are similar (Carrarini, Jackson, & Pin, 2009). If honest people care about other people's honesty and about not profiting from others' dishonesty, they should prefer to interact with honest partners. In such a case, we should find that honest people are more likely to leave their partners when the partner is dishonest than when the partner is honest (Hypothesis 2a).

The second possibility is that honest people care about being honest but not so much about their partner's honesty or about avoiding profiting from other people's dishonesty. If this is true, honest people will refrain from lying but will tolerate others' lies. Indeed, people care about seeing themselves as moral but at the same time seek to maximize personal profit (Mazar, Amir, & Ariely, 2008). To maintain such an honest self-image, people use various self-serving justifications when violating rules (Ayal & Gino, 2011; Shalvi, Gino, Barkan, & Ayal, 2015). People may thus also use self-serving justifications to financially profit from others' lies. For honest people, a dishonest partner may provide the best of both worlds—such a partner allows honest people to maintain their self-image, and at the same time profit from their partner's dishonesty. Such behavior is comparable with free-riding in a public-goods dilemma, in which people may profit from others' contributions to a public good while not contributing themselves (Andreoni, 1988). Here, the “public good” is generated by violating a rule (i.e., lying), and people can profit financially from others' lies while refraining from lying themselves. We label this behavior *ethical free-riding*, which we define as intentionally benefiting from other people's rule-violating behavior without violating rules oneself. If honest people are indeed willing to tolerate others' rule violations, we should find that honest people are more likely to leave their partners when the partner is honest versus than when the partner is dishonest (Hypothesis 2b).

Does partnering with a liar turn the honest dishonest?

In collaborative settings in which people work together, one's behavior may affect the partner's behavior. Indeed, unethical behavior, such as helping someone else cheat on an

exam, can be learned from observing peers, colleagues, and even strangers (O'Fallon & Butterfield, 2012). Merely seeing someone lie facilitates imitation of such behavior (Kocher, Schudy, & Spantig, 2017; Soraperra et al., 2017), being exposed to norm violations increases subsequent norm violation (Keizer, Lindenberg, & Steg, 2008), and the more prevalent corruption is in a country, the more likely people from that country are to violate rules (Gächter & Schulz, 2016). Furthermore, because people do not like to be excluded (Kurzban & Leary, 2001), and are even willing to violate rules that benefit the group, to avoid being excluded (Thau, Derfler-Rozin, Pitesa, Mitchell, & Pillutla, 2015), honest people might adapt their behavior and start lying to prevent their dishonest partners from leaving them.

We thus assessed the impact that one's dishonesty has on one's partner's dishonesty, that is, the contagiousness of dishonesty. To interpret how contagious dishonesty is when people can select their partners, we compared the partner-selection setting with two benchmark settings in which people cannot select a partner: one in which people are forced to interact with the same partner, and another in which people are forced to change partners. The degree to which dishonesty is contagious depends on the partner preferences of honest people (Hypotheses 2a and 2b). If honest people attempt to engage in ethical free-riding but are being abandoned by their dishonest partners, they may start lying to compensate for the monetary loss incurred by failed ethical free-riding attempts. Hence, we had no ex-ante predictions concerning how one's dishonesty would affect the partner's dishonesty.

Finally, we explored the overall financial impact of corrupt collaboration. We did so by comparing the three settings in terms of both the proportion of honest dyads and how effective peoples' lies were.

Method

Experimental setup and key variables. The key variable of interest was people's decision to switch versus stay with their partners. We study this decision by implementing a novel experimental setting in which participants were assigned a partner and could increase their payoff by coordinated lying. In the experimental Choice condition, participants could choose whether to stay with or switch their partner. To investigate how the freedom to choose a partner affects the contagiousness of dishonesty and the overall financial impact of coordinated lies, we compared the Choice condition with two additional conditions: a Forced Stay condition in which participants had to stay with their partner throughout the experiment, and a Forced Switch condition in which participants were forced to repeatedly switch their partner.

We employed a modified version of the dyadic die-rolling task (Weisel & Shalvi, 2015) in which participants are assigned a partner and are asked to report the outcome of a die roll presented on a computer screen (Kocher et al., 2017). Participants earn extra money if they and their partner report the same outcome, with higher outcomes corresponding to higher pay. Employing experimental tasks such as reporting die-roll outcomes to assess dishonesty is a common practice both in the lab and the field (Fischbacher & Föllmi-Heusi, 2013; for a meta-analysis, see Abeler et al., 2019). The task has good external validity. Lying in these tasks is correlated with diverse unethical behaviors such as not paying for public transport (Dai, Galeotti, & Villeval, 2018), being absent from work without a reason (Hanna & Wang, 2017), and not returning undeserved pay (Potters & Stoop, 2016).

Participants. Participants were recruited from the subject pool of the Center for Research in Experimental Economics and Political Decision Making (CREED) at the University of Amsterdam. In total, 372 participants (186 females, $M_{\text{age}} = 22.00$, $SD_{\text{age}} = 3.38$) took part in 19 experimental sessions. All sessions lasted approximately one hour and were run by two experimenters who answered participants' questions, if there were any. All data and materials are available on the Open Science Framework¹¹. Screenshots of the instructions also appear in the Supplemental Online Material (SOM)¹². No participants were excluded from any of the analyses, and all independent variables and manipulations are reported in the Method section and the SOM. Average earnings per participant were €18.68.

The dyadic die-rolling task. Participants were seated in individual cubicles in front of a computer screen, where they read the task instructions. Each participant was paired with a partner and assigned the role of first or second mover. Participants' roles as first and second movers remained fixed throughout the experiment. To avoid reputation concerns, we made sure that participants were not aware, during or after the task, with whom they were partnered. In the first stage of each round, the first mover observed a random die roll on a computer screen, presented as a video segment (Kocher et al., 2017), and was instructed to report the outcome. In the second stage, the second mover saw the outcome the first mover reported on the computer screen, and then observed an independent, random die roll and was asked to report it (see Figure 4.1A). Dyad members then received a summary of their own and their partner's report and a payoff summary for that round. After each block of three rounds,

¹¹ https://osf.io/eym9j/?view_only=09abbe27e68e4dc69b9b6bc6e14073dc

¹² <https://tinyurl.com/yyeu8y5q>

participants saw a block summary. Overall, the task lasted for 30 rounds (10 blocks), but participants were not aware of the total number of rounds in the task. Each participant had to correctly answer three questions regarding the task's rules to ensure task comprehension before starting the task (see the Supplemental Material). Per CREED lab rules, no experimental deception was used and decisions had real monetary consequences.

Payoff scheme. For every round, each participant received a fixed payment of €3 regardless of the outcomes reported. Beyond that amount, if the first and second mover reported the same outcome (i.e., a double), they received an extra payment associated with their reports. Specifically, each participant earned half of the reported outcome's worth in euros (e.g., if both reported "two," each earned an extra €1; if both reported "six," each earned an extra €3). If the dyad did not report a double, they received no extra payment (see Figure 4.1B). Thus, both participants had an incentive to misreport their die-roll outcomes and coordinate their reports. The first mover had an incentive to report the highest die-roll outcome possible, whereas the second mover had an incentive to match the first mover's report. Note that compared with a dishonest first mover, a dishonest second mover has more impact on the dyad's payoff, because the second mover can guarantee a double in every trial, whereas the first mover can only increase the potential value of the double.

At the end of the task, one block of three consecutive rounds was randomly selected for pay. Thus, on top of the fixed payment of €9 (3 rounds × €3) per block, participants earned a variable amount depending on their own and their partner's honesty, or lack thereof. The expected value if both dyad members report honestly was €0.29 ($[\text{€}0.5 + \text{€}1 + \text{€}1.5 + \text{€}2 + \text{€}2.5 + \text{€}3] \times 1/6 \times 1/6$) per round. On average, each member of a completely honest dyad would thus earn an extra payment of €0.88 (3 trials × €0.29) per block. Maximizing profit by reporting dishonestly, however, can increase participants' extra payment tenfold. In a fully dishonest dyad, the first mover would report the highest die-roll outcome (a six), corresponding to a payoff of €3, and a dishonest second mover would guarantee a double by matching this number. A fully dishonest dyad could thus earn an extra payment of €3 per round, leading to a total payment of €9 (3 trials × €3) per block. Importantly, we chose the payoff scheme so that by reporting honestly, participants would earn the usual hourly payment at our lab.

Because the die-roll outcomes were presented on the computer screen, we were able to identify for each report whether a participant was honest or not. Participants were informed that every time a double was reported in which one or both partners misreported the number

they observed, the dyad's earnings from the double would be subtracted from a planned donation to a charity (see SOM for details). Thus, the instructions made clear that it would be possible to determine whether participants misreported the die-roll outcome when data collection was completed.

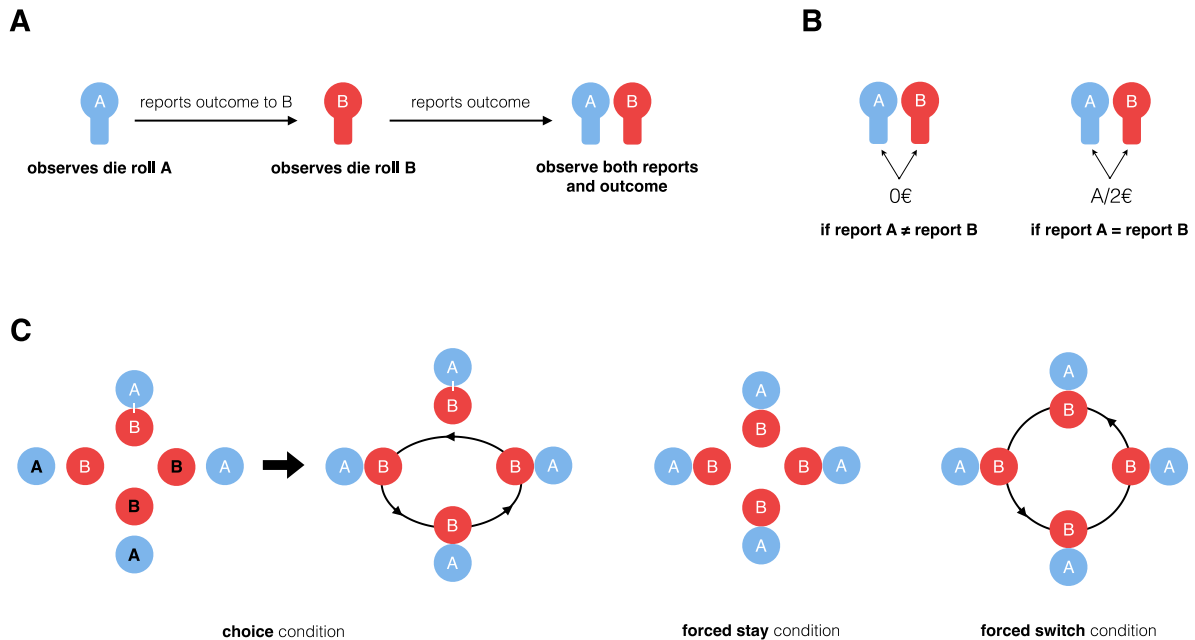


Figure 4.1. Task procedure, payoff scheme, and conditions. As shown in (A), the first mover, A, observes a die roll on a computer screen and reports the outcome to the second mover, B. B then observes an independent die roll on the computer screen and reports the outcome. Then A and B learn about both reports and the corresponding payoff. As shown in (B), if the number that B reports matches the number A reports (a double), each earns half ($A/2$) of the reported number in euros. Otherwise, they earn no extra pay. In the choice condition, participants indicated whether they wanted to switch partners (black letters in panel C) or stay with the same partner (white letters in panel C) after each block of three consecutive rounds. If at least one partner in a dyad wanted to switch, the dyad was split, and both members of that dyad were assigned a new partner for the next block. In the forced-stay condition, participants engaged in the task for 30 rounds with the same partner. In the forced-switch condition, participants were assigned a new partner after every block.

Partner selection. In the experimental Choice condition, participants were asked whether they wanted to stay with their partner or switch to a new partner after each block of three rounds. Participants could thus end their current partnership and look for a new partner. If at least one of the dyad members wanted to switch, and given that at least one more person from another dyad in the session wanted to switch, the dyad would split and both dyad members were reassigned a new partner (Figure 4.1C). For details on the algorithm used to

reassign participants, see SOM. We further ran two control conditions in which we did not allow for partner choice. In the Forced Stay condition, the first and second movers interacted with one another across all 30 rounds. In the Forced Switch condition, the first and second movers were forced to switch to a new partner after each block of three rounds. Participants read all the instructions before starting the experiment and received clarifications if they asked for them.

Analysis approach. In both the Choice and Forced Switch conditions, participants could potentially interact with all other first and second movers in their session. The nested nature of the Choice and Forced Switch conditions meant each session of 20 participants was treated as one independent observation. In the Forced Stay condition, each participant interacted with only one participant, making the dyad (rather than the session) one independent observation. Given the size of our subject pool, we opted to split our sample in a way that allowed us to obtain the most independent observations for each condition. Thus, in the Choice ($n = 160$) and in the Forced Switch ($n = 160$) conditions, we ran 8 sessions of 20 participants (=10 dyads) per session, and in the Forced Stay condition, we ran 3 sessions with a total of 52 participants (= 26 dyads), leading to a total sample of 372 participants. In total, we had 8 strictly independent observations in the Choice and Forced Switch conditions, and 26 independent observations in the Forced Stay condition. In previous work assessing the effect of partner selection on various economic games (e.g., Abbink, 2004; Page, Putterman, & Unel, 2005), data were collected from 144 to 256 participants, and there were between 6 and 9 independent observations for the partner-selection conditions. Our sample size compares favorably with these studies. We thus feel confident that our sample size and number of independent observations were sufficient to detect effects. All in all, our design allowed us to analyze both choice-based switching patterns and the contagiousness of dishonesty between participants and dyads operating within a large group and across a relatively long time horizon (30 rounds).

We performed data analysis on strictly independent observations, by aggregating across sessions (in the Choice and Forced Switch condition) or dyads (in the Forced Stay condition), or by fitting mixed-effects regression models (also called multilevel regression). To account for the nested structure (in dyads and participants), we specified two random intercepts, one for each participant (with multiple observations across rounds) and one for each unique dyad in these models. When the dependent variable was binary, we fitted a logistic mixed-effects regression model. Further, when the dependent variable had natural

boundaries (e.g., die-roll reports were restricted to be between one and six), we treated the data as left- and right-censored, to avoid biased estimates resulting from heteroscedasticity. All mixed-effects models were fitted with R and JAGS using a Bayesian approach with noninformative priors. We hence do not report p-values for these models, but instead report the Bayesian 95% confidence interval (see SOM for further details).

Results

We first focused on partner-selection decisions in the Choice condition, to test Hypotheses 1 and 2. We then compared the contagiousness of dishonesty and the overall financial impact of corrupt collaboration among the Choice, Forced Stay, and Forced Switch conditions.

Finding a partner in crime. In the Choice condition, we categorized each participant as an honest person (H) if he or she reported the outcomes that appeared on the computer screen in all three rounds in a given block, or as a liar (L) otherwise. This procedure led to four possible dyad types: LL (lying first mover, lying second mover; accounted for 50% of the dyads in our sample), LH (lying first mover, honest second mover; 4% of the sample), HL (honest first mover, lying second mover; 31% of the sample), and HH (honest first mover, honest second mover; 15% of the sample). Figure 4.2 depicts the proportion of participants who asked to switch partners on the basis of their role (first or second mover), separately for each dyad type. Because data were clustered in dyads and participants (across blocks) and the dependent variable was binary (0 = stay with partner, 1 = switch partner), we fitted a random-intercept logistic regression model to the data, as described above. Results revealed that dishonest people indeed sought a partner in crime, supporting Hypothesis 1. Dishonest first movers were much more likely to ask to switch partners when interacting with an honest versus a dishonest second mover, whereas dishonest second movers were much more likely to ask to switch partners when interacting with an honest versus a dishonest first mover (Figure 4.2, random-intercept logistic regression, HL estimate vs. LL baseline: $b = 3.52$, 95% confidence interval, or CI = [1.98, 5.27], LH estimate vs. LL baseline: $b = 6.53$, 95% CI = [4.55, 8.84]). Likewise, when both dyad members were dishonest, the dyad was very stable and its members asked to switch only 3.5% of the time (1.4% for first movers, 5.6% for second movers).

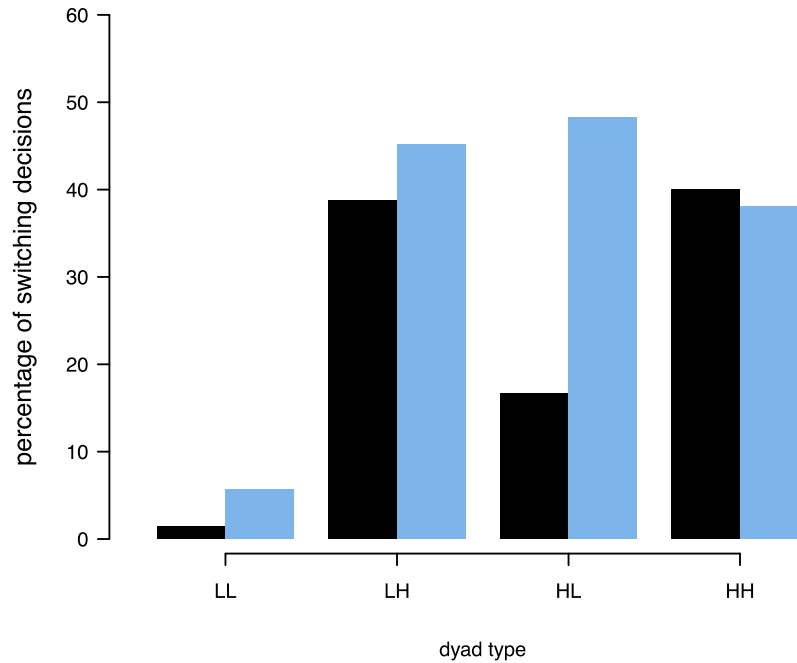


Figure 4.2. Proportion of first movers (black) and second movers (blue) who chose to switch partners, separately for each dyad type (LL = lying first mover, lying second mover; LH = lying first mover, honest second mover; HL = honest first mover, lying second mover; HH = honest first mover, honest second mover).

Furthermore, the more times a second mover matched a dishonest first mover's report (i.e., reported a double), the less likely the dishonest first mover was to ask to switch (Figure 4.3A, random-intercept logistic regression, number of doubles estimate: $b = -8.61$, 95% CI = $[-12.84, -4.83]$), and the more times a first mover reported the highest possible outcome, six, the less likely a dishonest second mover was to ask to switch (Figure 4.3B, random-intercept logistic regression, number of sixes estimate: $b = -1.30$, 95% CI = $[-1.62, -0.99]$). Hence, in line with Hypothesis 1, results showed that dishonest participants preferred partners who helped them secure the highest profit possible.

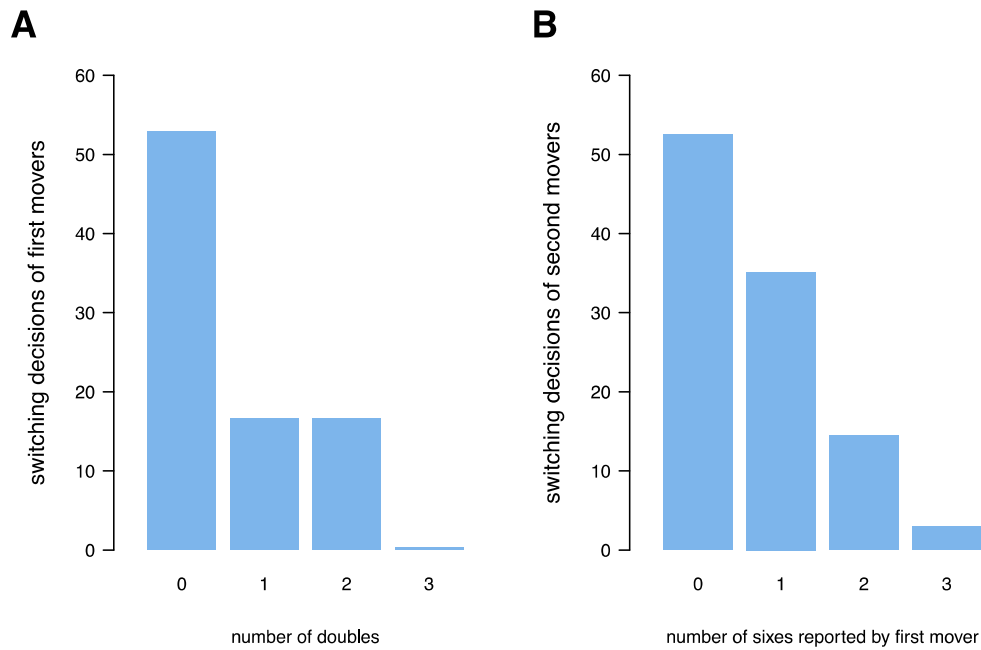


Figure 4.3. Proportion of dishonest first movers who chose to switch partners as a function of the number of doubles reported in the previous block (A) and proportion of dishonest second movers who chose to switch partners as a function of the number of sixes reported by the first mover in the previous block (B).

A preference for honest partners?

As predicted, people who were willing to violate rules to maximize profit systematically switched partners until they were paired with a like-minded, dishonest partner. But what did honest people do? Consider an honest first mover who is paired with a dishonest second mover. The first mover reports the number appearing on the screen, including the less profitable ones and twos, whereas the dishonest second mover matches every time. The honest first mover clearly prevents the dyad from maximizing potential profit, but at the same time, financially benefits from the second mover's dishonesty. In such a case, would the honest first mover switch in the hope of finding a like-minded honest partner? Or, alternatively, stay with the lying partner and enjoy both worlds—being honest while profiting from the other's dishonesty?

Results revealed that honest first movers systematically engaged in ethical free-riding, supporting Hypothesis 2b over Hypothesis 2a. Honest first movers were less likely to ask to switch when their partner (the second mover) was a liar versus an honest person (Figure 4.2, Table 4.1, line 4, random-intercept logistic regression, HL estimate vs. HH baseline: $b = -3.08$, 95% CI = [-4.37, -1.83]).

One may argue that an honest person facing a dishonest partner may not switch, because of the ambiguity associated with the partner's actions. It is not trivial for first movers to identify whether their partner was lying or not, because there was a chance that the partner actually observed the same outcomes as they did. Indeed, the likelihood of a second mover matching exactly one of the first mover's three die rolls equals 35%. If ambiguity, rather than ethical free-riding, is driving honest first movers' decision, we should observe that the more doubles a second mover reports, the less likely the first mover is to stay with this person. The reason is that with every additional double reported, the ambiguity about the second mover's dishonesty fades away. For example, the likelihood that a second mover would honestly match the first mover's reports three out of three times is less than 0.5%.

Our data suggest that honest first movers were ethically free-riding rather than refraining from switching a partner because of ambiguity concerns. As Figure 4.4A shows, the more times a second mover matched an honest first mover's report (i.e., reported a double), the less likely the honest first mover was to choose to switch (Figure 4.4A, random-intercept logistic regression, number of doubles estimate: $b = -3.36$, 95% CI = [-4.58, -2.27]). In particular, honest first movers chose to switch in 46% of the cases, if their partner did not match their reported numbers (possibly seeking a more profitable relationship), but only in 11% of the cases if their partner matched all of their reports, indicating they preferred a dishonest over an honest partner. Therefore, honest people seemed comfortable with profiting from the unambiguous dishonesty of their partner while they themselves remained honest.

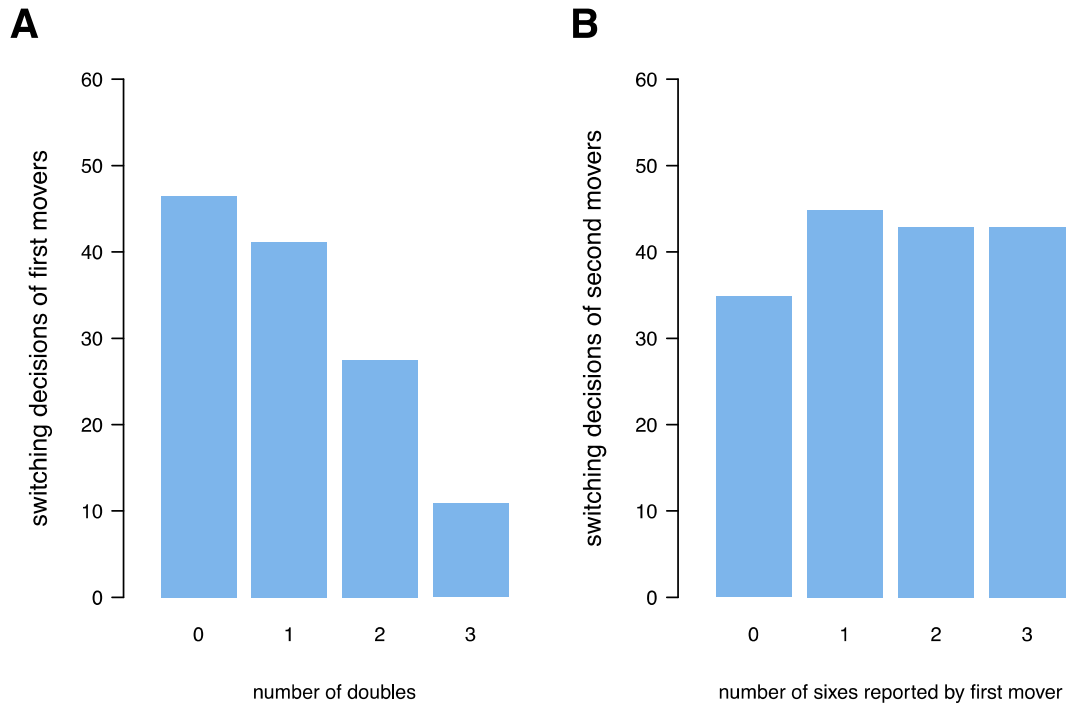


Figure 4.4. Proportion of honest first movers who chose to switch partners as a function of the number of doubles reported in the block (A) and proportion of honest second movers who chose to switch partners as a function of the number of sixes reported by the first mover in the block (B).

Table 4.1. *Switching decisions in the Choice condition.*

	Estimate	SD	95% CI
Intercept (first mover in HH dyad)	0.07	0.583	[-1.043, 1.239]
LL dyad	-6.58	0.966	[-8.512, -4.772]
LH dyad	-0.05	0.895	[-1.849, 1.682]
HL dyad	-3.08	0.645	[-4.368, -1.835]
Second mover	-0.07	0.748	[-1.494, 1.448]
Second mover × LL dyad	2.51	1.036	[0.537, 4.609]
Second mover × LH dyad	0.19	1.079	[-1.926, 2.289]
Second mover × HL dyad	3.37	0.821	[1.722, 4.970]
random intercept error term (dyad)	2.09	0.319	[1.513, 2.747]
random intercept error term (subject)	2.27	0.338	[1.643, 2.945]

Note. Results of the Logistic Random-Intercept Regression Modeling the Choice to Switch (0 = Stay, 1 = Switch) in the Choice Condition as a Function of Dyadic Relationship. We compared the interaction model with a null model and a model with only main effects to test whether the full model with interaction terms explained significantly more variance than a simpler model. Compared with the null model, a model with main effects significantly increased the likelihood of the data given the model, $\chi^2(4) = 149.39$, $p < .001$. Furthermore, a model with main and interactions terms significantly increased the likelihood of the data given the model, $\chi^2(3) = 21.24$, $p < .001$. CI = confidence interval; HH = honest first mover, honest second mover; HL = honest first mover, lying second mover; LH = lying first mover, honest second mover; LL, lying first mover, lying second mover.

The temptation for honest people to engage in ethical free-riding, however, depends on the mover's position. Compared with second movers (who can secure a double every time), first movers have no control over the reporting of a double and can determine only the doubles' potential worth (i.e., moving from an expected value of 3.5 if honest, to 6 if dishonest). As a result, for honest participants, having a dishonest second mover as a partner is more profitable than having a dishonest first mover as a partner. The temptation to ethically free ride for an honest first mover is hence higher than the temptation for an honest second mover. Specifically, the expected payment for an honest first mover increases from €0.87 ($€1.75 \times 1/6 \times 3$ rounds) per block to €5.25 ($€1.75 \times 1 \times 3$ rounds) per block when switching from an honest to a dishonest second mover. Thus, honest first movers face a temptation of a 503% premium for switching from an honest to a dishonest partner. The expected payment for

an honest second mover, by contrast, increases from €0.87 per block when paired with an honest first mover, to €1.50 ($€3 \times 1/6 \times 3$ rounds) per block when paired with a fully dishonest first mover that always reports sixes. Thus, honest second movers face a temptation of only a 72% premium for switching from an honest to a dishonest partner.

In line with the notably lower temptation for second movers to switch to a dishonest partner, we found no evidence for ethical free-riding among honest second movers. Honest second movers were as likely to choose to switch when they were paired with a dishonest or an honest first mover (Figure 4.2, Table 4.1, line 7, random-intercept logistic regression, second mover LH vs. HH estimate: $b = 0.19$, 95% CI = [-1.93, 2.29]). Furthermore, honest second movers were indifferent to the number of sixes reported by the first mover when they decided whether to switch or not (Figure 4.4B, random intercept logistic regression, number of sixes estimated: $b = -2.70$, 95% CI = [-10.69, 4.22]). Note that whereas honest first and second movers in our experiment may not have consciously calculated the exact expected value when choosing between switching or staying with their partners, their switching decisions were consistent with the notion that a dishonest second mover has more influence on the dyads' payoff than a dishonest first mover.

Taken together, supporting Hypothesis 1, dishonest people seek a dishonest partner who will allow them to secure high payoffs on the basis of mutual rule violations. Supporting Hypothesis 2b, honest people do not seek like-minded honest peers. Instead, honest first movers exploit the freedom to choose a partner and switch away from honest second movers. When paired with a dishonest second mover, many participants engaged in ethical free-riding. Ethical free-riding allows honest first movers to remain honest but secure high profits from the dishonest behavior of their partners.

The prevalence of ethical free-riders. We assessed the prevalence of ethical free-riding by classifying participants according to the most common choice strategy they exhibited throughout the experiment (in the Choice condition). In particular, we classified each participant's behavior on a block-by-block basis and then assigned unique labels for each individual depending on the behavior that was most prevalent for this person. Specifically, (a) ethical free-riders were honest people (i.e., reported the outcomes observed in all three rounds) who chose to switch when paired with an honest partner and stay when paired with a liar in most of the blocks; (b) Kantian truth-tellers were honest people who chose to switch when paired with a liar and stay when paired with an honest partner; (c) brazen liars were liars (i.e., misreported the die-roll outcome

in at least one out of three rounds) who chose to switch when paired with an honest partner and stay when paired with a liar; (d) tolerant liars were liars who chose to stay when paired with an honest partner; and (e) confused liars were liars who chose to switch when paired with a liar.

In total, most participants (70.6%) were classified as liars, with the large majority being brazen (89.4% brazen liars vs. 10.6% tolerant liars). Among the percentage of honest people (29.4%), the majority were ethical free-riders (63.9%) and a minority were Kantian truth-tellers (36.1%). The rank order of this type frequency was stable across different classification procedures (see SOM).

Participants were rather consistent in their types in the first and second half of the experiment. Cramér's ϕ captures the association between categorical (or nominal) variables, with 0 indicating a complete lack of association (i.e., lack of consistency) and 1 indicating a perfect association (i.e., full consistency). Participants' classification into one of the aforementioned five categories in the first half (i.e., the first four out of nine switching decisions) correlated with this person's classification in the second half of switching decisions (i.e., the subsequent five switching decisions), $\phi = 0.68$. This finding suggests that participants were rather consistent in their behavior, but some did modify their behavior (and thus type). We next assessed the extent to which partner's honesty and dishonesty drives such behavior modifications.

Contagiousness of dishonesty. Thus far, we have treated honesty as a stable individual difference and assessed behavior per block. For some people, however, the preference for honesty may not be stable (Gibson, Tanner, & Wagner, 2013) but may be influenced by the partner's behavior (Weisel & Shalvi, 2015). The Choice condition may critically affect the degree to which people influence one another, because in this setup, honest people face the possibility of having their dishonest partners leave them resulting in loss of potential earnings. We thus assessed the extent to which dishonesty is contagious in the Choice condition and compared it with the Forced Stay and Forced Switch control conditions.

To an extent, we found that lying is contagious. Participants were more likely to lie when their partner lied compared with when their partner reported honestly in the previous report (i.e., previous round for first movers and previous stage for second movers; Wilcoxon signed-rank test, $p < .01$). In the Forced Stay condition, participants' likelihood of lying in a given round *increased* by 39% when their partner previously lied versus reported honestly

(Table S2 in SOM, random-intercept logistic regression, partner's Time t-1 behavior estimate: $b = 0.33$, 95% CI = [0.02, 0.68]). The ability to choose partners did not reduce the contagiousness of dishonesty (Table S2, random-intercept logistic regression, Choice \times partner's Time t-1 behavior estimate: $b = -0.01$, 95% CI = [-0.42, 0.39]). Specifically, in the Choice condition, participants' likelihood of lying *increased* by 38% when their partner previously lied versus reported honestly. By contrast, compared to the Forced Stay condition, forcing people to switch partners reduced the likelihood that a participant would lie following a partner's lie (Table S2, random-intercept logistic regression, Forced Switch \times partner's Time t-1 behavior estimate: $b = -0.41$, 95% CI = [-0.81, -0.03]). In the Forced Switch condition, participants' likelihood of lying *decreased* by 8% after their partner had lied versus reported honestly, suggesting that forcing people to repeatedly switch partners can effectively curb the contagiousness of dishonesty.

These results show that, to a degree, dishonesty spread and some honest people were influenced by their dishonest partners and started lying as well. Across conditions, the contagiousness of dishonesty was stronger the longer the partners were together. In the Forced Stay condition participants interacted with their partner for 30 rounds, whereas in the Forced Switch condition, they interacted with each partner for only 3 rounds. In the Choice condition, some people choose to stay with their partners while other choose to switch, leading on average to longer relationships than in the Forced Switch condition and shorter relationships than in the Forced Stay condition. We found dishonesty to be the most contagious in the Forced Stay and the least contagious in the Forced Switch condition, suggesting that longer relationships increased the contagiousness of dishonesty.

Overall financial impact. The overall financial impact of corruption is composed of (a) the proportion of fully honest dyads and (b) the financial damage liars cause (i.e., their effectiveness). Interestingly, these two factors may (or may not) cancel each other out. A setting with a high prevalence of honest dyads (and thus less liars) in which lies are efficient (i.e., have a large financial impact) can produce an overall financial damage similar to a setting with a low prevalence of honest dyads (and thus more liars) in which lies are inefficient (i.e., have a small financial impact).

First, the overall magnitude of lying was similar to previously reported findings (Weisel & Shalvi, 2015; see SOM). In the Forced Stay condition, only 4.6% of dyads were fully honest (i.e., both the first and second mover were honest in a given block). The prevalence of honest dyads roughly doubled under Forced Switch (9.9%; Table S3 in SOM,

censored regression, $b = 0.17$, $SE = 0.07$, $p = .01$), and tripled under the Choice condition (14.5%; Table S3, censored regression, Choice condition estimate: $b = 0.22$, $SE = 0.07$, $p < .01$). Giving people the choice to switch partners led to a relatively high level of fully honest dyads.

Second, financial damage caused by liars is reflected in how effective lies are. Effective lies are those in which the first mover reports a high number and the second mover matches it. When both dyad members are coordinated on high numbers, lying secures high payoffs. When a dyad is poorly coordinated, some lies are not reciprocated, resulting in ineffective dishonesty. To assess the effectiveness of each dyad's lies, we assessed the proportion of brazen lies – those reporting a double of sixes, leading to the maximum pay of €3 per dyad member. As can be seen, in Figure 4.5, 31% of lies yielded the maximum financial damage in the Choice condition. This finding was similar to that in the Forced Stay condition in which 38% of lies resulted in the highest possible monetary return, $\chi^2(1, N = 178) = 0.48$, $p = .49$. By contrast, only 9% of lies resulted in maximum pay in the Forced Switch condition, which was lower than in the Forced Stay and Choice conditions, $\chi^2(1, N = 311) = 22.1$, $p < .01$, Bonferroni corrected for multiple comparisons. Analyzing the average effectiveness of lies across conditions led to the same conclusions (see SOM).

These findings suggest that high levels of coordinated dishonesty emerge not only from long-term relationships (Forced Stay), but also from dishonest people seeking to interact with one another (Choice). The Choice condition allowed people to find like-minded dishonest partners who would help them coordinate on lies and increase payoff. Forcing people to switch and thus restricting their ability to coordinate reduces the effectiveness of their lies.

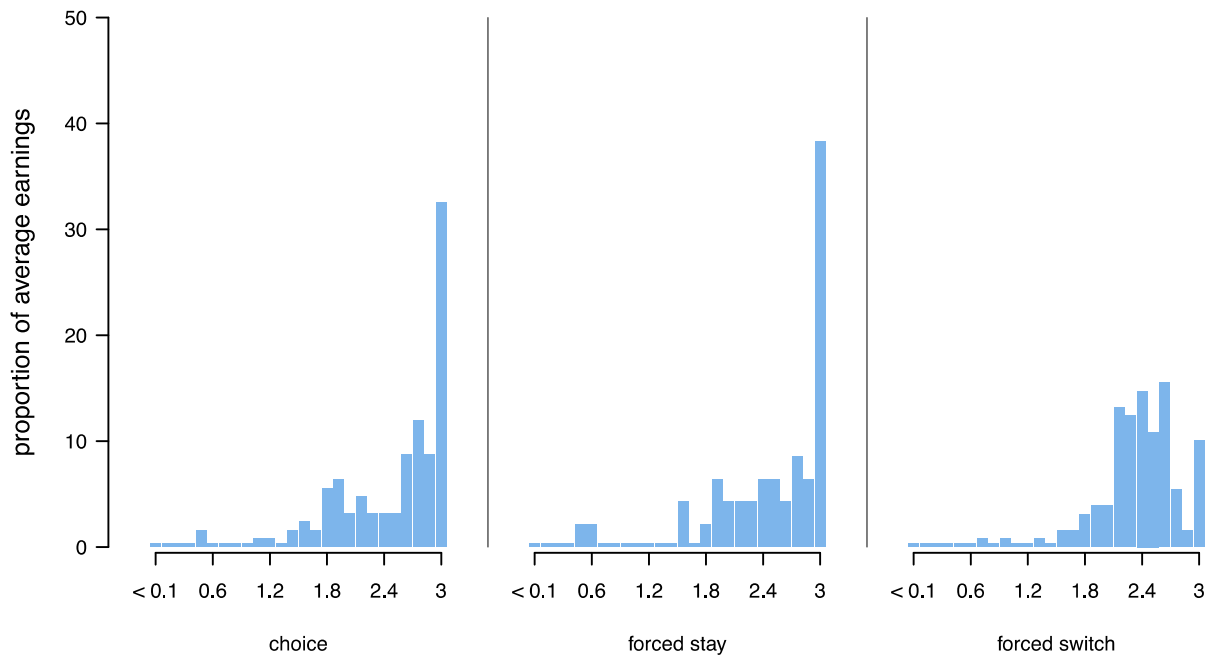


Figure 4.5. Distribution of average earnings per round (in Euros) of dishonest participants, separately for each condition.

Taken together, whereas the Choice condition contained a relatively large proportion of fully honest dyads (and thus fewer liars), participants who lied were very efficient in causing financial damage. By contrast, the Forced Switch condition contained a relatively lower proportion of fully honest dyads (and thus more liars); however, those who lied were less efficient, causing less financial damage. Seen in this light, the finding that the accumulated effects of corrupt collaboration (i.e., money earned by misreporting) did not vary greatly between the Choice and Forced Switch conditions is not surprising, because the two factors canceled each other out. Overall, the average earnings from dishonesty (per dyad, per block) was somewhat lower in both the Choice ($M = 10.79$, $SD = 6.83$; 10.67% reduction) and Forced Switch ($M = 10.23$, $SD = 6.65$; 15.31% reduction) conditions, compared with the Forced Stay condition ($M = 12.08$, $SD = 5.94$). These differences did not reach statistical significance (Mann-Whitney U- test, $p = .17$).

Discussion

Human societies are based on cooperation. People, however, occasionally engage in joint acts of rule violations, aimed at maximizing personal gains, ultimately hurting societal welfare. Understanding how such corrupt collaborations emerge and spread is important for both theoretical and practical reasons. Here, we investigated how honest and dishonest people

choose their partners and how the ability (vs. inability) to choose a partner affects the spread and overall impact of corrupt collaboration. Results show that to secure high payoffs, dishonest people prefer interacting with dishonest partners. Interestingly, honest people also exploit the possibility of being able to choose their partners. The majority of the honest participants preferred to interact with dishonest partners, a financially beneficial but morally dubious choice. Those honest people engaged in ethical free-riding—they remained honest and “kept their hands clean” but at the same time benefitted from the fact that their partners did the “dirty work” and lied. Choosing one’s partner enabled participants to indirectly violate rules by opting for the “naive approach” and exploit the moral wiggle room (Dana, Weber, & Kuang, 2007). Those people had evidence to support the fact that they were honest, and they could have claimed they were unaware of any misreporting taking place. In retrospect, they may admit that perhaps things were too good to be true but may still cling to their honesty as a path to resist punishment and condemnation.

Beyond partner-selection choices, we further explored how the ability to select a partner affects the contagiousness of dishonesty and the financial impact of coordinated lies. By comparing the Choice condition with two control settings in which people were forced to stay or switch partners, we found that honesty is malleable and lies are contagious, especially when forcing partners to stay together. Forcing people to switch partners helps counteract the contagiousness of lies. But such a policy does not reduce the overall financial impact of corrupt collaboration. We found a relatively high proportion of fully honest dyads in the Choice condition, compared with the Forced Switch condition. Conversely, liars were more efficient in the Choice condition, leading to similar levels of the overall financial impact of corruption across conditions.

Our results relate to work showing that the ability to select one’s group members increases cooperation in public-goods dilemmas. Indeed, working with like-minded people increases cooperation (Gächter & Thöni, 2005), and being able to select group members facilitates cooperation (Gross, Méder, Okamoto-Barth, & Riedl, 2016; Page et al., 2005). Our findings demonstrate, however, that allowing people to choose their partners, and thereby increasing cooperation, may carry negative consequences. In settings in which cooperation can be achieved by joint rule violations, providing the freedom to choose partners may be detrimental. Seen in this light, cooperation in and of itself is neutral—it carries neither good nor bad externalities. The externalities that cooperation has on society (increased welfare in

some cases, corruption in others) are what make high levels of cooperation, and by extension, the freedom to select partners, desirable or undesirable.

A key observation in our work is the existence of ethical free-riders – people who intentionally benefit from others’ rule-violating behavior without violating rules themselves. This observation resonates with work on passive cheating (also known as lying by omission), showing that individuals refrain from correcting mistakes, made by a computer or another person, as long as those mistakes are beneficial to them (Pittarello, Rubaltelli, & Motro, 2016). Whereas ethical free-riding and passive cheating are similar in terms of profiting from not taking an action, ethical free-riding goes beyond tolerating mistakes (caused by nature or another person) and indicates that people will also tolerate others’ intentional rule violations.

Studying how people ignore others’ intentional rule violations seems promising for future research. For instance, whereas participants in our experiment had minimal reputation concerns (i.e., interacted anonymously), in real life, people’s reputation is often at stake. Assessing how ethical free-riding is shaped by reputation concerns, whether people engaging in such behavior are negatively (or perhaps positively) evaluated by their peers, and how ethical free-riders behave when they cannot rely on other people to violate rules for their own profit seem worthy of further investigation. Additionally, motivating the honest to act is a key goal of research on whistle-blowing programs. Whistle blowing requires people to act on the warning signals they observe regarding others’ intentional rule violations, and to forego potential benefits from tolerating such misconducts (Waytz, Dungan, & Young, 2013). To create successful interventions aimed at curbing corruption and encouraging whistle blowing, we need to further uncover the underlying mechanisms that make people turn a blind eye to others’ rule violations and engage in ethical free-riding.

An interesting aspect to consider is the goal that corrupt collaboration serves, namely, whether it is driven by a motivation to avoid potential losses or secure potential gains. Going back to our opening example, Volkswagen’s employees might have manipulated the engine’s software not only because of a desire to secure a bonus, but also for fear of losing their jobs (Goodman, 2015). Indeed, the motivation to approach a positive outcome facilitates dishonesty (Pulfrey & Butera, 2013; Schweitzer, Ordóñez, & Douma, 2004), but the motivation to avoid losing money may sometimes be an even stronger motivator to lie for profit (Grolleau, Kocher, & Sutan, 2016; Schindler & Pfattheicher, 2017). Future research exploring how approach versus avoidance motivation shapes the emergence and spread of corrupt collaboration and ethical free-riding seems promising.

Finally, the results reported here have organizational implications. Employees can often choose their collaborators, units, or work groups on the basis of their organization's job-rotation policies. Rotating facilitates innovation (Coşgel & Miceli, 1999) and increases employees' satisfaction (Ho, Chang, Shih, & Liang, 2009). Interestingly, some organizations (e.g., the United Nations) employ forced-rotation policies, as mimicked by our Forced Switch condition. Others provide employees with the freedom to choose whether or not to rotate (voluntary rotation in, e.g., the European Commission; Fontaine & Tang, 2006), as mimicked by our Choice condition. As we show here, the freedom to rotate comes with a moral hazard. Providing the possibility to choose with whom to interact can increase the spreading and contagiousness of corrupt collaboration and enables ethical free-riding. The possible advantages of voluntary job-rotation policies should be carefully weighed with the possible risks, especially in environments with high collusion and corruption potential.

Conclusion

The freedom to select partners is important for the establishment of trust and cooperation. As we show here, however, it is also associated with potential moral hazards. For individuals who seek to keep the risk of collusion low, policies providing the freedom to choose one's partners, should be implemented with caution. Relying on people's honesty may not always be sufficient, because honest people may be willing to tolerate others' rule violations if they stand to profit from them. Our results clarify yet again that people who are not willing to turn a blind eye and stand up to corruption should receive all praise.