Being biased against friends to appear unbiased

Article *in* Journal of Experimental Social Psychology · July 2018

DOI: 10.1016/j.jesp.2018.05.009

3 authors:

Alex Shaw
University of Chicago
32 PUBLICATIONS **688** CITATIONS
SEE PROFILE

Shoham Choshen-Hillel
Hebrew University of Jerusalem
15 PUBLICATIONS **158** CITATIONS
SEE PROFILE

Eugene M Caruso
University of Chicago
40 PUBLICATIONS **985** CITATIONS
SEE PROFILE

Journal of Experimental Social Psychology 78 (2018) 104–115

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

Being biased against friends to appear unbiased

Alex Shaw[a,*], Shoham Choshen-Hillel[b], Eugene M. Caruso[c]

[a] University of Chicago, USA
[b] The Hebrew University of Jerusalem, Israel
[c] University of Chicago, Booth School of Business, USA

ARTICLE INFO

Handling editor: Robbie Sutton

ABSTRACT

In contexts where fairness is important, people attempt to avoid the appearance of partiality. Although such efforts to avoid appearing partial can often reduce biases, we argue that, at times, such efforts can actually lead people to be biased *against* their friends. We theorize that people do so because they recognize that benefitting their friends may be viewed by others as partial. This argument makes two key predictions, which we investigated in eight studies using workplace scenarios. First, we predicted and found that, when the decision was public, allocators were reluctant to give a bonus to a deserving employee when that employee was a friend rather than a non-friend. In private, however, participants were willing to give the bonus to the deserving person whether she was a friend or a non-friend, suggesting that their public behavior was aimed at avoiding the appearance of bias. Second, we predicted and found that allocators' reluctance to give a bonus to a deserving friend is mediated by their beliefs that others would find this behavior to be unfair. We discuss the theoretical and practical implications of this bias resulting from a desire to avoid appearing partial.

1. Introduction

A manager has to give a bonus to one of two equally deserving employees, one of whom is her friend. A judge in a high school debating competition has to decide which of two finalists to vote for, one of whom is a student from his alma mater. A coach has to decide which of two players should start in the championship game, one of whom is the coach's niece. It is well known that people show favoritism toward those close to them, such as family, friends, or in-group members. However, in this paper we argue that, in public situations where people are concerned with avoiding the appearance of partiality, people may sometimes actually be biased *against* their friends. Indeed, if the decision will be sufficiently scrutinized by others, decision-makers may be reluctant to give a benefit to a friend, even if that friend deserves it, because they may be worried that giving the benefit will be perceived as partial. We start by reviewing previous research that highlights the delicate balance that people must strike between a desire to favor their friends and a desire to avoid the appearance of partiality. We then investigate the phenomenon of being biased to appear unbiased in eight studies.

1.1. Favoritism and fairness

Research in psychology and economics has demonstrated time and again that people show preferential treatment to their friends and to those who have treated them favorably in the past (Cox, 2004; Delton, Krasnow, Cosmides, & Tooby, 2011; DeScioli & Kurzban, 2009a; Fiske, 1992; Gurven, 2004, 2006; Rand, Arbesman, & Christakis, 2011; Shaw, DeScioli, & Olson, 2012; Trivers, 1971). In the context of friendships, people uniquely value their friends and preferentially offer support to them (Clark & Mills, 1979; Tooby & Cosmides, 1996). People also like friends who offer them preferential treatment (Barakzai & Shaw, in press; Cole & Teboul, 2004; Kenny, Mohr, & Levesque, 2001; Laurenceau, Barrett, & Pietromonaco, 1998), and who preferentially take their side during disputes (DeScioli & Kurzban, 2009a, 2011; DeScioli, Kurzban, Koch, & Liben-Nowell, 2011; Pietraszewski, 2016). Indeed, failing to take a friend's side—by remaining neutral in a dispute, for instance—can damage a friendship (Shaw, DeScioli, Barakzai, & Kurzban, 2017). Further, there is extensive evidence that people are biased in favor of their in-group (Balliet, Wu, & De Dreu, 2014; Brewer & Silver, 1978; Tajfel, Billig, Bundy, & Flament, 1971). There is also rampant evidence of favoritism and cronyism in behavioral economic games, in the workplace, and in politics (Abbink, Irlenbusch, & Renner, 2002; Barr & Serra, 2009; Brick, Palmon, & Wald, 2006; Dungan, Waytz, & Young, 2014; Reinsch & Gardner, 2014). Taken together, these findings suggest that people are drawn toward favoritism—giving more based on friendships, reciprocal relationships, and shared group identity.

* Corresponding author.
E-mail address: ashaw1@uchicago.edu (A. Shaw).

https://doi.org/10.1016/j.jesp.2018.05.009
Received 3 August 2017; Received in revised form 23 May 2018; Accepted 30 May 2018

Although people are undoubtedly biased toward their friends, their wish to show favoritism toward their allies sometimes conflicts with their fairness concerns (DeScioli & Kurzban, 2009b, 2013; Dungan, Waytz, & Young, 2014; Niemi, Wasserman, & Young, 2018; Shaw, 2016; Waytz, Dungan, & Young, 2013). Here we develop the idea that people's concern with fairness may not only de-bias individuals by reducing their favoritism toward friends, but that in some circumstances it may cause them to be biased against their friends.

It has long been known that people care about fairness and respond negatively to inequity—when they or others receive unequal resources for equal work (Adams, 1965; Damon, 1977; Dawes, Fowler, Johnson, McElreath, & Smirnov, 2007; Messick, 1995; Shaw & Choshen-Hillel, 2017; Sweeney & McFarlin, 2004). This drive toward fairness, and people's desire to avoid creating inequity among others, can even lead people to incur personal costs to reduce inequality, or to waste resources rather than distribute them inequitably (Choshen-Hillel & Yaniv, 2011; Gordon-Hecker, Rosensaft-Eshel, Pittarello, Shalvi, & Bereby-Meyer, 2017; Mitchell, Tetlock, Mellers, & Ordonez, 1993; Shaw & Knobe, 2013).

In explaining these findings, a recent partiality account of fairness argues that people's fairness concerns are driven by their desire to avoid the appearance of partiality rather than inequality or inequity per se; inequality is fair if it is created in an impartial manner, but it is unfair if it is created in a partial manner (Choshen-Hillel, Shaw, & Caruso, 2015; Shaw & Knobe, 2013; Tyler, 2000). Partiality is defined as showing favoritism toward others based on their individual identity (DeScioli & Kurzban, 2013). Equality is often viewed as impartial, as it does not force the allocator to favor one recipient over another (Choshen-Hillel, Shaw, & Caruso, in press; Messick, 1995). Yet, inequality may also be viewed as impartial, to the extent that the distribution is based on a socially agreed upon justification (e.g., merit, need, or a randomization procedure; see Hook & Cook, 1979; Shaw, 2016; Shaw & Olson, 2012, 2014; Thibaut & Walker, 1975; Tyler, 2000). For example, if Josh does a better job than David and is being paid more for his work, this distribution is unequal but may still be perceived as impartial because it is based on merit (i.e., doing a better job). Merit is a culturally agreed upon justification that makes the resulting inequality impartial. Importantly, unequal pay for equal work can also be considered fair if it is created using an impartial procedure like a coin flip (Choshen-Hillel et al., 2015, in press; Tyler, 2000). There is now extensive evidence that people will create inequity, and believe it is fair to do so, if it is based on an appropriate justification (Choshen-Hillel, Shaw, & Caruso, 2015; Shaw, 2013; Shaw, Choshen-Hillel, & Caruso, 2016).

The partiality account further suggests that people's fairness concerns are focused on avoiding the appearance of partiality. Specifically, it suggests that allocators anticipate that showing favoritism toward one party will cause the disadvantaged parties, as well as uninvolved third parties, to condemn them for their decision; they therefore try to avoid making partial allocations, particularly in circumstances where others can observe their allocations (e.g., in public; Andreoni & Bernheim, 2009; Batson, Kobrynowicz, Dinnerstein, Kampf, & Wilson, 1997; Leary, Allen, & Terry, 2011; Levitt & List, 2007; Piazza & Bering, 2008; Reis & Gruzen, 1976; Shaw et al., 2014).

Here we build on the partiality account of fairness to argue that the desire to avoid the appearance of partiality can cause people to endorse a different kind of biased decision-making, namely, bias against their friends. In the same way that there are factors that make inequality appear less partial (e.g., merit, randomized procedures), there are factors that make inequality appear more partial. One particularly strong cue that may increase the likelihood of perceiving partiality is a pre-existing alliance of some kind, such as a friendship. Because people frequently favor their friends, observers may jump to the (often correct) conclusion that an allocator who gives a friend more than someone else is showing favoritism, and is therefore being unfair. Thus, we expect that if the recipient who benefits from an allocation is the allocator's friend (vs. not), people will be more likely to assume that the decision-maker's choice was based on partiality. Supporting this idea, we know that when an allocator gives more to one person than another for no good reason (e.g., in the absence of any difference in merit), people sometimes infer that the allocator is friends with the beneficiary (Kleiman-Weiner, Shaw, & Tenenbaum, 2017; Liberman & Shaw, 2017).

We build on this reasoning to suggest that when an allocator creates an unequal yet equitable allocation (i.e., paying a friend more than someone else based on merit), observers may tend to attribute the inequality to the friendship and infer partiality. As a result, a concern with being condemned for partiality may cause an allocator to be reluctant to give more to a friend even when there is a merit-based justification for doing so. Thus, although avoiding the appearance of partiality will often lead to less bias by forcing individuals to counteract their tendency to favor their own interests and those of their allies, we predict that, in some predictable circumstances, the desire to avoid appearing partial can lead people to be biased *against* their friends. That is, they will systematically treat their friends in a worse manner than they treat their non-friends.

1.2. The current studies

In eight studies using workplace scenarios, we explored the predictions of the partiality account. In these scenarios, allocators were asked to decide whether to give a larger bonus to one recipient who did a slightly better job or to treat the two recipients equally (e.g., flipping a coin to decide who gets the bonus). We predicted that when the decision is public, people would be more likely to give the bonus to the slightly more deserving person when this person is merely a colleague than when the colleague is also a friend. We made this prediction because we thought participants might worry that others would judge them as unfair (or partial) if they gave the bonus to their friend. In contrast, when the recipient is not a friend, they would not fear they would be judged as unfair if they base their decision in merit. Further, we predicted that the tendency not to give a bonus to a slightly more deserving friend would be less prominent in private than in public because this tendency is driven by people's concern with the outward appearance of bias, and not with actual bias. Finding evidence for this pattern of results would provide crucial support for the idea that people choose to be biased against their friends to avoid the appearance of partiality. In line with these predictions about allocators, we also predicted that third-party observers would judge an allocator as more unfair when she gives more to a deserving friend than to a deserving colleague who was not her friend (Study 3a and 3b). We further hypothesized that allocators would anticipate this fact, and that this anticipation would mediate their biased decisions outlined above (Study 4 and 7). We close by investigating some options that can help allocators avoid this biased decision-making, and by considering contexts in which people might not be concerned with appearing partial (Study 6).

2. Study 1

In Study 1, we investigated a unique prediction of the partiality account of fairness: that allocators will be biased against their friends in public by being reluctant to create inequality that favors a friend, even when the friend merits a larger reward. To investigate this possibility, we asked participants to imagine that they were managers at a firm that had to allocate a bonus between two employees. Participants were told that one of the employees received slightly better evaluations from others in the office (i.e., the difference in merit was established by others, not the manager herself). Participants learned that the two employees had each been given a bonus and that they (as managers) had two options for allocating an additional bonus: they could give the extra bonus to the employee who received slightly higher evaluations, or they could give it to neither employee.

We manipulated whether the manager's decision was private or

public and whether the more highly evaluated employee was the manager's friend or a mere colleague (i.e., someone with whom the manager does not have any personal relationship). When the manager's decision was private, we predicted that managers would generally be willing to allocate the additional bonus to the employee who did a slightly better job, whether the recipient was a friend or a mere colleague. This prediction is in line with previous work that suggests people favor giving more to others based on merit, such that better performance warrants higher compensation (e.g., Hook & Cook, 1979). When the decision was public, however, we predicted that the manager's decision would depend on whether the recipient was a friend or a mere colleague. When the recipient was a mere colleague, we predicted that allocators would be quite willing to give more to the recipient who did the better job because there is no strong reason for anyone to suspect this decision to be the result of favoritism (and if they do not give it to anyone, it would be wasteful). However, when the recipient was a friend, we predicted that allocators would be more reluctant to give the bonus in public because they will be worried that it may appear to be predicated on friendship rather than merit. Specifically, we predicted that one cell of our design—when the decision was public and the allocator was giving to a friend—would be different than the other three cells, with participants being less likely to give the resource in this case. Alternatively, if people are simply interested in doing what they think is fair or impartial without concern for how others might interpret their choice, then they should be willing to give the bonus to the more deserving employee whether the decision is public or private and whether the recipient is a friend or a mere colleague.

2.1. Method

2.1.1. Participants

We recruited 218 participants on Amazon Mechanical Turk (63% female, M_{age} = 30.68 years, SD_{age} = 9.87). For all studies, participation was restricted to participants from the United States, and compensation was 25 cents for completing the study, which took about 5 min. Sample size was determined before data collection started. We aimed for ~55 participants per cell. All measures and manipulations are reported in the paper. We included a comprehension check in each study to prevent the possibility of gaining a reputation on MTurk for not including comprehension checks, but we decided a priori to include all participants in our studies, regardless of this question.

2.1.2. Procedure

The task was presented online using Qualtrics survey software. Participants read instructions explaining that they would read a vignette and be asked to make a choice about how to distribute resources. These vignettes were adapted from previous work (Choshen-Hillel, Shaw, & Caruso, 2015; Shaw & Knobe, 2013). The study used a 2(Relationship: Friend, Mere Colleague) × 2(Observation: Public, Private) between-participants design. In the friend conditions, participants read:

> Imagine that you are working at a large company. You have been asked to decide how to assign a bonus to two employees, Mark and Dan. They both currently make the same amount each year, do the same job, and have received almost identical evaluations. Mark and Dan received evaluations from three other employees.
>
> Mark had a slightly more positive evaluation from 2 of the 3 raters and Dan had a slightly more positive evaluation from 1 of the 3 raters. Mark is a really close friend of yours and everyone in the company knows he is a really close friend of yours.
>
> The company can give a total of a $500 bonus, but based on institutional rules must assign the bonus in hundred dollar increments. Mark and Dan have each been given a $200 bonus. What would you do with the other $100 bonus?

Participants in the public condition were then told that "Everyone will know you made this decision," whereas participants in the private condition were told instead that "No one will know you've made this decision." Then, all participants were asked to choose whether to "Give the $100 bonus to Mark" or "Give the $100 bonus to neither." The mere colleague condition was identical except for the sentence "Mark is a really close friend of yours and everyone in the company knows he is a really close friend of yours," which was replaced with "You do not know either employee personally." Note that we specified that the evaluations came from other raters, not the allocator herself. We did this to minimize the possibility that participants might try to correct for their own biased evaluations. After participants made their decision, they were asked to fill in some brief demographic information.

2.2. Results

We conducted a logistic regression analysis with relationship (friend or mere colleague) and observation (private or public) as independent variables and decision as the dependent variable. We also included an interaction variable, computed by multiplying the two centered variables. The analysis revealed a marginal main effect of relationship, Wald(1, N = 218) = 2.96, p = .085, Exp(B) = 1.96, such that participants were somewhat more likely to give the bonus when the recipient was a mere colleague (91 out of 104, 87.5%) than a friend (85 out of 114, 75%). There was also a main effect of observation, Wald(1, N = 218) = 4.23, p = .040, Exp(B) = 2.23, such that participants were more likely to give the bonus in private (97 out of 110, 88%) than in public (79 out of 108, 73%). Importantly, these main effects were qualified by a significant interaction between relationship and observation, Wald(1, N = 218) = 4.70, p = .030, Exp(B) = 5.45. See Fig. 1.

We used planned comparisons (Fisher's exact test) to follow up on this interaction. In line with our hypothesis derived from the partiality account, we found that people were less likely to give the bonus to the slightly more deserving employee when the decision was public and the employee was their friend (36 out of 59, 61%) as compared to all other conditions: friend private (49 out of 55, 89%, p < .001, φ = 0.32), mere colleague private (48 out of 55, 87%; p = .003, φ = 0.30), and mere colleague public (43 out of 49, 88%; p = .002, φ = 0.30). None of the other comparisons approached significance, ps > .900.

2.3. Discussion

Consistent with the partiality account of fairness, we found that participants were less willing to publicly give an additional bonus to a deserving colleague when that colleague was a friend than not. Participants were only reluctant to give a bonus based on merit when they were giving the bonus to a friend in public. In private, participants

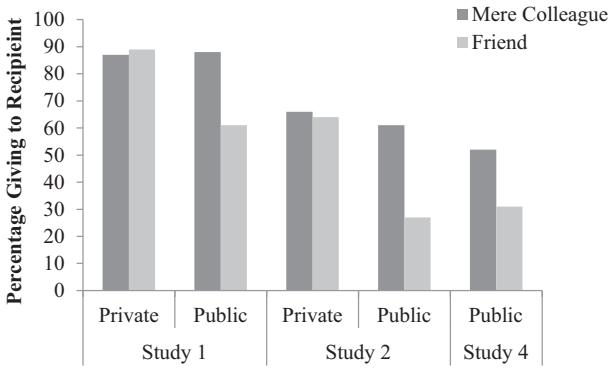

Fig. 1. Percentage of participants opting to give more to the recipient (who did a slightly better job) in public versus private and when the recipient was their friend versus a mere colleague in Study 1, Study 2, and Study 4. For Study 4, the numbers are from the "slight merit condition."

were equally willing to give a bonus based on merit to a friend or a colleague who was not their friend. The fact that participants were willing to give a bonus to their friend based on merit in private provides evidence against the possibility that participants were merely trying to override their internal bias in favor of their friend; if this were the case, participants should be more reluctant to give to a deserving friend in both public and private (further in a supplemental study we found that when the two recipients were equally deserving, people were biased in favor of their friend in private but not in public, see Study S1). The fact that we found this reluctance to give a bonus to a friend only when the decision was public suggests that participants were indeed trying to avoid the appearance of partiality and not simply trying to correct for a personal bias.

3. Study 2

The findings from Study 1 were consistent with the idea that people try to avoid the appearance of partiality. We found that people were reluctant to create merit-based inequality by giving more to a friend, but only when their decision was public. However, Study 1 involved a situation in which the only options were to give the bonus to one employee or give it to no one. Although such wasteful options have been used in past research (Choshen-Hillel, Shaw, & Caruso, 2015; Shaw & Olson, 2012), they may not be representative of situations in which allocators could likely find another way to assign the bonus if they did not want to use slight differences in merit.

Thus, in Study 2, we attempted to replicate our findings from Study 1 using the same basic design, but in a scenario in which participants could either give the bonus based on merit or flip a coin and give the bonus to the winner. If people agree that it is fair to give more to the person who did a slightly better job, then flipping a coin might be considered unfair because it risks the possibility that the less deserving employee will get more. Thus, just as in Study 1 and in line with the partiality account, we again expected that in three of the conditions people would favor giving the bonus to the more deserving recipient: when giving to a mere colleague or to a friend in private and when giving to a mere colleague in public. However, we predicted that participants would be reluctant to give the bonus to their friend in public, and so they would instead favor flipping the coin to avoid appearing partial. Again, if people are concerned with avoiding partiality rather than just its appearance, then we would expect similar behavior in all four conditions.

3.1. Method

3.1.1. Participants

We recruited 216 participants on MTurk (54% female, M_{age} = 36.42, SD_{age} = 12.22).

3.1.2. Procedure

The procedure was similar to Study 1, involving a sharing scenario between two employees (this time named David and Jonathan) in which David did a slightly better job than Jonathan. The study again used a 2(Relationship: Friend, Mere Colleague) × 2(Observation: Public, Private) between-participants design. Here, the choice given to participants was between giving the extra bonus to David or flipping a coin and giving the bonus to the winner. In the friend case, participants read the following:

> Imagine that you are a manager at a big company. You are about to assign bonuses to two of your employees, David and Jonathan. When you go over the employees' evaluations over the past year, you find that David and Jonathan both performed very well, and received nearly identical evaluations from their peers. David has been more productive – but not by much.
>
> The company can give a total of a $500 bonus, but based on institutional rules must assign the bonus in hundred dollar increments. David and Jonathan have each been given a $200 bonus. You now have to decide what to do with the other $100 bonus. You can either flip a coin and if it lands on heads you'll give the $100 bonus to David and if it lands on tails you'll give the $100 bonus to Jonathan. Or you can just give the $100 bonus to David.
>
> Note that everyone in the company knows that David is a friend of yours and that you like him.

Participants in the public condition were then told, “Everyone in the company will eventually find out what your decision was.” In the private condition, participants were told instead, “No one in the company will find out what your decision was.” The mere colleague condition was similar, except that the sentence “Note that everyone in the company knows that David is a friend of yours and that you like him” was not included. Finally, in all conditions, participants were asked, “How would you assign the extra bonus?” and could decide to give the bonus to David or to flip a coin and give the bonus to the winner. Participants then filled in some brief demographic information.

3.2. Results

We conducted a logistic regression analysis with relationship (friend or mere colleague) and observation (private or public) as independent variables and choice as a dependent variable. We also included an interaction variable, computed by multiplying the two centered variables. The analysis revealed a main effect of relationship, Wald(1, N = 216) = 6.73, p = .009, Exp(B) = 2.12, such that participants were more likely to give the bonus to a mere colleague (73 out of 115, 63%) than to a friend (46 out of 101, 46%). There was also a main effect of observation, Wald(1, N = 216) = 8.86, p = .003, Exp(B) = 2.37, such that participants were more likely to give the bonus in private (70 out of 108, 65%) than in public (49 out of 108, 45%). These main effects were qualified by a significant interaction between relationship and observation, Wald(1, N = 216) = 5.59, p = .018, Exp(B) = 3.93. See Fig. 1.

We used planned comparisons (Fisher's exact test) to examine the relevant comparisons. In line with our hypothesis, we found that people were less likely to give more to a recipient when it was their friend in public (14 out of 51, 27%) as compared to all other conditions: friend private (32 out of 50, 64%, p < .001, φ = 0.37), mere colleague private (38 out of 58, 66%; p < .001, φ = 0.38), and mere colleague public (35 out of 57, 61%; p < .001, φ = 0.34). None of the other comparisons approached significance, ps > .840.

3.3. Discussion

In a scenario that gave participants the option of using a random device to avoid wasting resources, we replicated our results from Study 1: Participants were reluctant to give an additional bonus to their friend in public, even when that friend was slightly more deserving of the bonus. We did not observe this pattern when giving to a friend in private or when giving to a colleague in public or private, suggesting that participants were indeed reluctant to reward their friend in public because they were concerned that their decision to give to a friend might be seen as partial. Note however that we would not expect people to be biased against their friends in public in cases where it was clear to everyone that their friend was more deserving than the other person (e.g., when their friend's performance was far superior), which is indeed what we found in a supplemental study (see Study S2).

4. Study 3a

Studies 1 and 2 found that allocators were biased against their friend in public by being less willing to give them a bonus than to a mere colleague; however, this result does not determine whether

observers actually judge such allocators as more unfair for giving to the friend. In Studies 3a and 3b, we examine whether third-party evaluations are aligned with the behavior of allocators in Studies 1 and 2; that is, do people evaluate giving more to a friend based on merit as more unfair than giving more to a mere colleague based on the same merit?

We theorized that the allocators' decisions in Studies 1 and 2 were partly driven by allocators' assumption that others would condemn them as unfair for giving a bonus to their friend, and thus they avoided doing so. However, recent work has found that people often mispredict how much people care about fairness in such allocation settings (Cooney, Gilbert, & Wilson, 2016). That is, decision-makers could falsely assume that third parties will judge them as unfair, when in fact they do not. Thus, an alternative possibility is that third parties' evaluations would not correspond to the decisions of allocators.

To investigate this question, in Study 3a, we presented participants with vignettes very similar to those in Study 1 and asked them to evaluate how fair the different decisions were. Specifically, participants were told that a manager could give a bonus to an employee who did slightly better than another employee or they could give the bonus to neither employee. What varied between conditions is what the allocator decided to do (i.e., to give the bonus or not) and whether the recipient in question was a friend or a mere colleague of the allocator.

We predicted a correspondence between people's choices in the previous studies and third party evaluations, such that third parties would evaluate the decision to give more to a friend as more unfair than the decision to give more to a non-friend colleague because the pre-existing relationship would act as a cue that the allocation might be based on partiality.

4.1. Method

4.1.1. Participants

We recruited 191 participants on MTurk (60% female, M_{age} = 30.71 years, SD_{age} = 9.64).

4.1.2. Procedure

Participants were randomly assigned to read about a manager who was faced with a bonus allocation between two employees. The manager's decision was very similar to the one made in the public conditions from Study 1. Participants were told that one employee had done slightly better than the other, and that the manager could either give the bonus to that employee or to neither one. Participants were then told that the manager decided to give the bonus to the employee who did slightly better or to give it to neither employee. The employee who did slightly better was either a friend of the manager or a mere colleague. The study thus used a 2(Relationship: Friend, Mere Colleague) × 2(Choice: Give, Not Give) between-participants design. Participants were asked to make a rating of their agreement or disagreement with the statement, “What the manager did was fair” on a scale that ranged from 1 (*Strongly Disagree*) to 7 (*Strongly Agree*). Participants were asked brief demographic information.

4.2. Results

A 2(Relationship: Friend, Mere Colleague) × 2(Choice: Give, Not Give) ANOVA on people's fairness evaluations revealed no main effect of relationship, $F(1, 187) = 0.01$, $p = .928$, $\eta_p^2 = 0.00$. There was a marginal main effect of choice, $F(1, 187) = 3.45$, $p = .065$, $\eta_p^2 = 0.02$. Importantly, however, there was a significant relationship by choice interaction, $F(1, 187) = 17.74$, $p < .001$, $\eta_p^2 = 0.087$ (see Fig. 2). We followed up on the interaction with planned comparisons. We found that giving to the recipient was seen as more fair when giving to a mere colleague ($M = 5.54$, $SD = 1.20$) than to a friend ($M = 4.57$, $SD = 1.59$), $t(92) = 3.42$, $p < .001$, $d = 0.69$. Relatedly, giving to no one was seen as more fair when the potential recipient was a friend ($M = 5.13$, $SD = 1.81$) than a mere colleague ($M = 4.11$, $SD = 1.85$), $t(90) = 2.66$, $p = .009$, $d = 0.56$. Further, we found that when the recipient was a mere colleague, participants rated giving the bonus as more fair than giving to no one, $t(92) = 4.50$, $p < .001$, $d = 0.92$. However, when the recipient was a friend, giving the bonus was not seen as more fair than giving the bonus to no one, $t(95) = 1.60$, $p = .113$, $d = 0.33$.

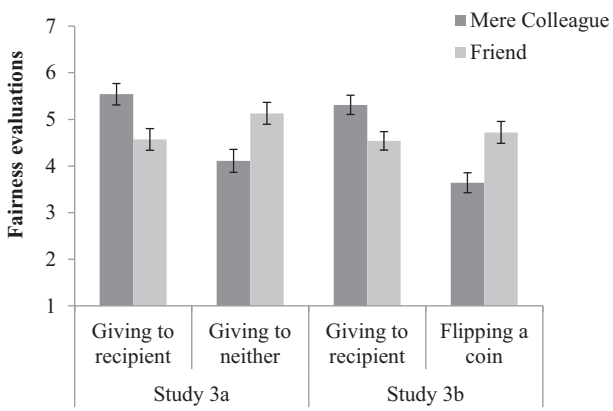

Fig. 2. Fairness evaluations in Studies 3a and 3b for a third-party's decision based on whether the recipient was their friend versus a mere colleague.

4.3. Discussion

In line with our prediction, we found that participants judged giving more to a deserving recipient as less fair when that recipient was a friend rather than a mere colleague. These results converge with the behavioral data from Study 1: Allocators in Study 1 were reluctant to give an additional bonus to their friend, and observers in Study 3a found it unfair for allocators to give this bonus to a friend. These results show that decision-makers seem to correctly infer that third parties (on average) will find it unfair for them to give more to a friend than to a non-friend.

5. Study 3b

Study 3b further examined whether people would judge a decision maker to be less fair if they chose to give a bonus to their friend than to a non-friend colleague. Here, we used a decision context similar to the one in Study 2, in which the decision maker's option was to give the bonus to the recipient who did a better job (friend or mere colleague) or to flip a coin to decide who should get the extra bonus. We predicted that people's fairness evaluations would correspond with the behavior observed in Study 2, in that observers would judge it as less fair to give the bonus to a deserving friend than a deserving colleague.

5.1. Method

5.1.1. Participants

We recruited 226 participants on MTurk (61% female, M_{age} = 33.95 years, SD_{age} = 10.91).

5.1.2. Procedure

Participants read about a manager who was faced with a decision very similar to the one made in the public conditions from Study 2. Specifically, the participants were told about a manager who had to decide whether to give a bonus to an employee who had done slightly better than another employee, or to give it according to a coin flip. Participants were then told that the manager decided to give the bonus the employee who did a slightly better job or to flip a coin and give it to the winner. The study thus used a 2(Relationship: Friend, Mere Colleague) × 2(Choice: Give, Flip Coin) between-participants design.

Participants then filled the same fairness measure from Study 3a and asked brief demographic information.

5.2. Results

A 2(Relationship: Friend, Mere Colleague) × 2(Choice: Give, Flip Coin) ANOVA on participants' fairness evaluations revealed no main effect of relationship, $F(1, 222) = 0.52$, $p = .471$, $\eta_p^2 = 0.002$. There was a main effect of choice, $F(1, 222) = 12.06$, $p < .001$, $\eta_p^2 = 0.05$, where giving the bonus ($M = 4.90$, $SD = 1.49$) was seen as more fair than flipping the coin ($M = 4.12$, $SD = 1.83$). Importantly, there was a significant relationship by choice interaction, $F(1, 222) = 18.92$, $p < .001$, $\eta_p^2 = 0.079$. See Fig. 2.

We followed up on this interaction with planned comparisons. We found that giving the bonus to the recipient was seen as more fair when the recipient was a mere colleague ($M = 5.31$, $SD = 1.37$) than a friend of the manager ($M = 4.53$, $SD = 1.51$), $t(123) = 3.00$, $p = .003$, $d = 0.54$. Additionally, flipping a coin was judged as more fair when giving to a friend ($M = 4.72$, $SD = 1.73$) than when giving to a colleague ($M = 3.64$, $SD = 1.77$), $t(99) = 3.09$, $p = .003$, $d = 0.62$. Further, when the recipient was a colleague, participants judged giving the bonus as more fair than flipping a coin, $t(112) = 5.66$, $p < .001$, $d = 1.06$. Yet, when the recipient was a friend, there was no significant difference in the fairness evaluations for giving the bonus versus flipping the coin, $t(110) = 0.61$, $p = .543$, $d = 0.12$.

5.3. Discussion

As in Study 3a, we found that participants judged giving more to a somewhat more deserving recipient as less fair when that recipient was a friend of the decision maker than merely his colleague. These results provide another demonstration that people do judge others more harshly for sharing unequally to favor friends, even when giving more to the friend is in line with merit.

6. Study 4

Studies 1–3 found that people were more reluctant to give a deserved bonus to a friend than a mere colleague (at least when the decision was public) and that third parties evaluated giving the bonus to be less fair when it was given to a friend than a mere colleague. In Study 4, we explore whether allocators anticipate that others will judge them as unfair for giving more to a deserving friend as compared to a mere colleague and if so, whether such anticipation mediates their choice.

6.1. Method

6.1.1. Participants

We recruited 100 participants on MTurk (49% female, $M_{age} = 37.69$, $SD_{age} = 12.17$).

6.1.2. Procedure

The procedure of Study 4 was an exact replication of the friend and mere colleague public conditions from Study 2, with the same vignette and relationship manipulation (friend or mere colleague). The only difference was that after the participants made their choice between giving the bonus to the harder-working employee and flipping a coin between the two employees, they were now asked to evaluate the fairness of these two choice options. Specifically, they were told "Having made that choice, please indicate the extent to which you agree or disagree with the following." Participants were then asked to state their agreement or disagreement with two statements presented in a counterbalanced order: "Other employees will think it was fair for me to give the extra bonus to David rather than flip the coin" and "Other employees will think it was fair for me to flip a coin rather than giving the extra bonus to David." The scale for each question ranged from 1 (*Strongly Disagree*) to 7 (*Strongly Agree*).

6.2. Results

6.2.1. Choice

A Fisher's exact test revealed that participants reported being less likely to give the bonus to a slightly more deserving employee when this employee was their friend (15 out of 48, 31%) as compared to when this employee was a mere colleague (27 out of 52, 52%), $p < .044$, $\varphi = 0.21$. This finding replicates the finding from Study 2.

6.2.2. Fairness

Participants anticipated that other people would find it less fair if they gave the bonus to the slightly more deserving employee when the employee was the their friend ($M = 2.81$, $SD = 1.38$) than when the employee was a mere colleague ($M = 4.48$, $SD = 1.49$), $t(98) = 5.80$, $p < .001$, $d = 1.17$. Relatedly, participants anticipated that other people would find it more fair if they flipped a coin when the employee was a friend of theirs ($M = 5.81$, $SD = 1.08$) than a mere colleague ($M = 4.79$, $SD = 1.76$), $t(98) = 3.46$, $p = .001$, $d = 0.70$.

6.2.3. Mediation analysis

Finally, we tested whether fairness (difference score between participant's evaluation of fairness for the "giving the bonus" and "flipping the coin" questions) mediated the effect of condition on decision. We found that condition (friend vs. mere colleague) influenced decision ($b = 0.86$, $p = .038$) as well as fairness ($b = -2.69$, $p < .001$). Reported fairness, in turn, influenced decision ($b = -0.86$, $p < .001$), and the inclusion of fairness in the analysis reduced the effect of condition on choice ($b = -0.92$, $p = .143$) suggesting full mediation. A bootstrap analysis (10,000 bootstrapped sample) revealed that the 95% bias-corrected confidence interval for the size of the indirect effect excluded zero (1.32, 4.02), suggesting a significant indirect effect of fairness (MacKinnon, Fairchild, & Fritz, 2007). Thus, as expected, fairness mediated the relationship between choice in the mere colleague and friend conditions.

6.3. Discussion

The results of Study 4 replicated our previous findings by showing that people report that they would be more reluctant to give a bonus to a slightly more deserving employee if this employee were their friend rather than a mere colleague. Study 4 extended these findings by demonstrating that the participants were more likely to report they thought that others would judge them as unfair if they gave the bonus to the slightly more deserving employee who is their friend rather than a mere colleague. Importantly, this judgment about the unfairness of the allocation predicted participants' likelihood of choosing to give the bonus to the slightly more deserving employee or flip the coin. It is worth noting that participants' prediction about others evaluations were fairly well aligned with the actual evaluations of third parties that were obtained in Study 3b. The results of Study 4 provide important support for our theoretical account by showing that people's tendency to be biased against their friends is driven by their anticipation that others will see them as unfair if they give more to their friends.

7. Study 5

In our previous studies, we investigated cases in which people could be biased against their friend by defaulting to an impartial strategy—such as an equal division or a coin toss. We deliberately chose to make the "biased against friend" option to be equal because we hypothesized that people might think that such options appear fair. We thought people would be less likely to be biased against their friend in cases where they had to explicitly give the bonus to a less deserving employee. Unlike an equal division or a coin toss, giving more to a less

deserving employee is not an easily defensible strategy and so we expected that allocators would be unlikely to choose this option. However, it is possible that people would take this option if it were given to them in order to appear especially unbiased. To explore this possibility, we tested again a case where a manager could give an extra bonus to their slightly more deserving friend or flip a coin, but we now added a third option of giving the bonus to the less deserving employee.

7.1. Method

7.1.1. Participants

We recruited 100 participants on MTurk (49% female, M_{age} = 39.50, SD_{age} = 11.30).

7.1.2. Procedure

We pre-registered our hypothesis and the planned sample size (https://aspredicted.org/gz4bk.pdf). The procedure was an exact replication of the public conditions from Study 2 with the same vignette and comprehension checks (i.e., we again manipulated if the recipient was a friend or a mere colleague). The only difference was that in Study 5 participants were given three options: give the bonus to the slightly more deserving employee (David), to the less deserving employee (Jonathan), or flip a coin.

7.2. Results

7.2.1. Choice measure

A 2 × 3 Fisher's exact test revealed that the distribution of responses differed between the two conditions, p = .010. Further, in line with our expectation, we found that people very infrequently chose to give to the less deserving employee: zero participants out of 49 in the friend condition and one participant out of 51 in the mere colleague condition chose to do so. We then conducted a 2 × 2 Fisher's exact test on the other two options (giving to the slightly more deserving worker or flipping a coin). In line with our previous results, we found that people were less likely to give more to the slightly more deserving employee when he was their friend (27 out of 49, 55%) as compared to when he was a mere colleague (40 out of 50, 80%), p = .010, φ = 0.27.

7.3. Discussion

Study 5 replicated our previous results, finding that participants were less willing to give a bonus to a slightly more deserving recipient when that recipient was their friend rather than a mere colleague. We further found that although participants would opt to use an impartial procedure rather than giving the bonus to their friend, they would not opt to give the bonus to a less deserving employee. That is, the desire to avoid the appearance of bias may drive people to pursue equal division or random procedures that will be ostensibly fair, but in this context it did not cause them to choose to give more to someone less deserving.

8. Study 6

Our previous studies have found that allocators are reluctant to give a bonus to a slightly more deserving friend in cases where they are concerned with appearing biased toward that friend. We also found that the allocators' concern is reasonable: third parties will indeed judge allocators as unfair for giving a resource to a slightly more deserving friend. This means that in these situations, the allocator is in what seems like a no-win dilemma: she can give the resource to the friend and appear unfair or she can flip a coin and not appear unfair, but then she is being biased against her friend. One potential option that an allocator could use to overcome this dilemma is to elect to have someone else make the decision—effectively recusing herself from the decision based on a lack of ability to be impartial in this situation (Frost, 2005). In Study 6, we explore if allocators would actually choose to let another manager make the decision. We hypothesize that allocators will be more likely to let another manager make the decision when one of the recipients is their friend, rather than when both of the recipients are mere colleagues, because in the friend case they are more likely to be concerned about being or appearing biased.

8.1. Method

8.1.1. Participants

We recruited 100 participants on MTurk (57% female, M_{age} = 34.41, SD_{age} = 10.63).

8.1.2. Procedure

We pre-registered our hypothesis and the planned sample size The procedure was an exact replication of the public conditions from Study 2 with the same vignette and manipulation of relationship (friend or mere colleague). The only difference was that participants were given an additional third choice option about the allocation of the extra bonus: they could let another manager make the decision. Thus, participants could either give the bonus to the employee who was slightly more deserving, flip a coin, or let another manager decide.

8.2. Results

8.2.1. Choice measure

A 2 × 3 Fisher's exact test revealed that the distribution of responses differed between the friend and mere colleague conditions, p < .001. We conducted follow up Fisher's exact tests. First, a planned comparison compared the number of participants who opted to let another manager decide rather than made the choice themselves. We found that participants were much more likely to let another manager decide when one of the potential recipients was their friend (23 out of 51, 45%) rather than a mere colleague (4 out of 49, 8%), p < .001, φ = 0.52. Second, we compared the rates of participants choosing to give the bonus or flip the coin among those who opted to make a choice (i.e., did not recuse themselves). Here, we found no significant difference in participants' tendency to give the bonus to a friend (15 out of 28, 53.5%) rather than a mere colleague (31 out of 45, 69%), p = .219, φ = 0.15, although this pattern was in the same direction as our previous studies.

8.3. Discussion

In line with our expectation, we found that participants very rarely recused themselves when they were allocating resources between two mere colleagues; it was only when one of the recipients was a friend that a substantial percentage (about half) of participants recused themselves. This seems consistent with the idea that participants understand the predicament that they are in when having to allocate resources to a recipient who is a friend and try to avoid this predicament by letting another manager decide. Of course, in some situations it might be difficult to find another qualified manager to do so, but these results suggest at least one way that people can overcome these dilemmas in which they may feel pressured to be biased against their friends in order to avoid appearing biased.

9. Study 7

Although impartiality is highly valued in many official roles, it is not always important; there are many contexts in which partiality should not be seen as unfair. Partiality may be very unfair in the office, politics, and the criminal justice system (Shaw, 2016). Yet partiality will likely be seen as less unfair in many day-to-day friendship and familial contexts (e.g., Fiske, 1992). Indeed, third parties tend to judge people negatively for *not* preferentially helping loved ones in some contexts (Hughes, 2017). Although people still may worry about

accusations of unfairness and partiality in their day-to-day interactions with friends outside of the workplace, such accusations will often be less frequent and carry less weight. As noted above, in many friendship contexts people expect favoritism (Shaw, Descioli, Barakzai, & Kurzban, 2017). Indeed, a completely impartial person would not make a good friend; friendship requires that someone puts your interests over others' interests (for discussion see Shaw, 2016).

Therefore, in Study 7, we compared participants' allocation behavior when allocating concert tickets that were either a workplace resource or a personal resource. As before, there were an uneven number of resources and so participants could either give the resource to the slightly more deserving person or flip a coin. We also asked participants to predict how others would evaluate their decision if they chose to give the resource or flipped a coin. In line with our previous studies, we predicted that when distributing a workplace resource, people would be reluctant to give their friend the additional resource and think that others would perceive doing so as unfair. However, when distributing their personal resources, allocators would be more likely to give the additional resource to their friend and also think that others would perceive it as fair. Indeed, we predicted that allocators' expectations about others' fairness evaluations would mediate their decisions.

9.1. Method

9.1.1. Participants

We recruited 101 participants on MTurk (49% female, M_{age} = 37.05, SD_{age} = 12.14).

9.1.2. Procedure

The procedure was similar to our previous studies. Participants were asked to imagine being a manager and had to allocate 5 tickets to the musical Hamilton. In both conditions, the potential recipient who could be given the extra ticket was their friend. In one condition, the tickets were an office resource and in the other condition, they were a personal resource. The office resource condition read as follows:

> Imagine that you work at a big company. At the company there are two employees who you know, David and Jonathan.
>
> The company recently gave you 5 tickets to the Broadway musical Hamilton to distribute as a bonus between the two coworkers.
>
> When you go over the employees' evaluations over the past year, you find that David and Jonathan both performed very well, and received nearly identical evaluations from their peers. David has been more productive – but not by much. David is also a close personal friend and everyone in the company is aware of this.
>
> You have given David and Jonathan 2 tickets each. You now have to decide what to do with the other ticket that the company gave you. You can either flip a coin and if it lands on heads you'll give the remaining ticket to David and if it lands on tails you'll give the ticket to Jonathan. Or you can just give the last ticket to David.
>
> Again, David is a close personal friend of yours and everyone in the company will eventually find out what your decision was.
>
> How would you assign the extra bonus?

The personal resource condition read as follows:

> Imagine that you work at a big company. At the company there are two employees who you know, David and Jonathan.
>
> Your wife recently gave you 5 tickets to the musical Hamilton as a personal birthday present. However, you have a business trip, so you will be unable to go. Your wife suggested giving them to someone else and you have decided to give your personal tickets to your coworkers. David is also a close personal friend and everyone in the company is aware of this.
>
> You have given David and Jonathan 2 tickets each. You now have to decide what to do with the other ticket that your wife gave you. You can either flip a coin and if it lands on heads you'll give the remaining ticket to David and if it lands on tails you'll give the ticket to Jonathan. Or you can just give the last ticket to David.
>
> Again, David is a close personal friend of yours and everyone in the company will eventually find out what your decision was.
>
> How would you assign the extra bonus?

Participants could either give the bonus to David or flip a coin. Note, we attempted to make these conditions very similar by making the participant be the manager of the recipients of the two potential recipients in both conditions. However, in the personal resource condition we did not mention that the friend did a slightly better job because we were worried this would make participants think about this even more as an office context, which might make them worried about appearing biased. Importantly, mentioning that the friend did a better job overall should, if anything, work against our predicted result—the fact that the friend did slightly better in the work context should have motivated people to give the bonus to him. Indeed, we suspect that an even larger number of participants would have flipped the coin in the office scenario if we made the two employees equally deserving. After reading one of the two vignettes, participants were given the same two choice options they had in most of the previous studies: give the resource to the friend or flip a coin. They also filled out the fairness measures from Study 2, rating how fair they thought others would evaluate the decision to give to the friend or flip a coin.

9.2. Results

9.2.1. Choice measure

In line with our hypothesis, a Fisher's exact test revealed that participants reported that they were less likely to give more to their friend when it was an office resource (14 out of 48, 29%) than when it was a personal resource (32 out of 53, 58.5%, $p < .003$, $\varphi = 0.31$).

9.2.2. Fairness

Participants predicted that others would think that giving the bonus to the recipient would be seen as less fair when the bonus was an office resource ($M = 2.69$, $SD = 1.42$) than when it was a personal resource ($M = 4.13$, $SD = 1.70$), $t(99) = 4.61$, $p < .001$, $d = 0.92$. Relatedly, participants predicted that others would think that flipping a coin would be seen as more fair when the resource was an office resource ($M = 5.71$, $SD = 1.38$) than a personal resource ($M = 5.09$, $SD = 1.62$), $t(99) = 2.04$, $p = .044$, $d = 0.41$.

9.2.3. Mediation analysis

We tested whether fairness (difference score between people's evaluation of fairness for the “giving the resource” and “flipping the coin” questions) mediated the effect of condition on choice. We found that condition (office vs. personal resource) influenced choice ($b = -1.23$, $p = .004$) as well as fairness ($b = 2.06$, $p < .001$). Reported fairness, in turn, influenced choice ($b = -0.62$, $p < .001$), and the inclusion of fairness in the analysis reduced the effect of condition on fairness evaluations ($b = -0.45$, $p = .392$) suggesting full mediation. A bootstrap analysis (10,000 bootstrapped sample) revealed that the 95% bias-corrected confidence interval for the size of the indirect effect excluded zero (-0.57, -2.34), suggesting a significant indirect effect of fairness (MacKinnon, Fairchild, & Fritz, 2007). Thus, as expected, fairness mediated the relationship between choice in the friend and mere colleague conditions.

9.3. Discussion

Study 7 replicated our previous findings by demonstrating that

people are biased against their friends in allocating office resources. Importantly, it extended our findings by demonstrating that people are not biased against their friends when dividing personal resources. This difference in choice behavior was mediated by participants' expectations about how others would judge their decision in terms of fairness. That is, participants thought that giving to their friend would be seen as less fair when distributing an office resource as compared to a personal resource and, correspondingly, were less likely to give the resource to their friend when it was an office resource as compared to a personal resource. These results demonstrate that participants are particularly likely to be biased against their friends when they are in an official role that requires them to be impartial. Interestingly, we still found that a number of participants opted to flip a coin even when sharing a personal resource. This may be because the decision we described still occurred within a work context in which the participant was asked to imagine being a manager of the two recipients. This provided a tight control, but it may have increased participants' desire to appear impartial even with a personal resource. Indeed, if the context were completely interpersonal among friends who did not work together, it is quite possible that we would have observed even lower rates of flipping the coin and that participants would be even more willing to favor their friend and think it is fair to do so.

10. General discussion

Eight studies provide converging support for the idea that people may be biased against their friends in order to avoid appearing partial. We found that although people were generally willing to use merit to give an additional bonus to a slightly more deserving employee, they were reluctant to do so when the deserving employee was their friend and their decision was public. Specifically, when their decision was observed by others, participants were less likely to give a bonus to a slightly more deserving employee when that employee was their friend than not (Studies 1, 2, 4, 5, and 7). When that same decision was private, participants were far more likely to give more to their slightly more deserving friend than when the decision was public, supporting our argument that allocators are trying to avoid the *appearance* of partiality rather than partiality itself (Studies 1 and 2). Consistent with allocators' behavior, observers judged allocators more harshly when they gave to a friend than to a non-friend, even when this was justified by slight differences in merit. Specifically, third parties inferred more unfairness when a person gave more to a slightly more deserving friend than to a non-friend (Studies 3a–3b). These expectations about other people's fairness evaluations mediated the effect of choice (Study 4). We further found many people chose to recuse themselves from the decision (by letting another manager decide) when one of the recipients was a personal friend, suggesting that participants recognize this tension between appearing unbiased and being unbiased (Study 6). Finally, we found that this pressure to avoid partiality is particularly acute in settings that officially require impartiality (i.e., when distributing office resources). That is, people were much more likely to favor their friend in public when dividing their personal resources rather than office resources (Study 7). Taken together, our results suggest that, at least in some contexts, people are motivated to avoid the appearance of partiality, and that this desire can drive them to be biased against their friends in publically observable circumstances.

10.1. Avoiding the appearance of partiality and resulting bias

The present results provide the strongest evidence to date that people's fairness concerns are driven by their desire to avoid the *appearance* of partiality. Although previous work has demonstrated that such fairness concerns are related to partiality (Choshen-Hillel, Shaw, & Caruso, 2015; Shaw & Olson, 2014), this work did not distinguish between a concern with appearing partial and concern with being partial. We have long known that people are more likely to be impartial and fair in public than private, suggesting they suppress partiality in public (e.g., Andreoni & Bernheim, 2009; Dana, Cain, & Dawes, 2006; Levitt & List, 2007; Reis & Gruzen, 1976). However, the current results suggest that not only are people more likely to be fair in public than in private, but that they are even willing to go against what they might privately think is fair in order to avoid the public appearance of bias. That is, they forsake *being* impartial to *appear* impartial. Further, in line with the partiality account, the current results reveal that people's behavior corresponds with third party evaluations: the cases where people were likely to infer unfairness based in partiality (i.e., where one was giving to a friend) were the same cases in which people were reluctant to create inequality based on merit. This pattern of results is consistent with the notion that people avoid the appearance of partiality to avoid accusations of unfairness from others.

Understanding how people navigate the dynamics between impartiality and favoritism is critically important because these tradeoffs permeate many aspects of people's lives. Whistle blowers, managers, politicians, judges, and even little league coaches must navigate these tradeoffs between fairness and loyalty (Dungan, Waytz, & Young, 2014; Niemi et al., 2018; Shaw, Descioli, & Olson, 2012; Waytz, Dungan, & Young, 2013). Although we examined this phenomenon in a specific context, the notion that avoiding partiality can lead one to be biased applies to a wide range of domains. For example, imagine that a psychology faculty member was recently promoted to a position of power within the administration. At first glance, this seems like an obvious win for psychology, as this professor is expected to look out for their ingroup (Balliet, Wu, & De Dreu, 2014; Brewer & Silver, 1978). However, as the results above make clear, the story need not be so simple. It is easy to imagine that the newly promoted administrator might sometimes second-guess their decisions because of concerns about appearing biased. That is, even if the psychology department deserved a new hiring line, the allocator might be slightly more reluctant to give it to them. Our results do not imply that such allocators must engage in biased decision-making (indeed, many of our participants still used merit even when they were concerned about appearing biased), but they do highlight a delicate pressure that people are sometimes faced with when they consider such decisions.

It is worth noting that the partiality account does not predict that people will never be impartial in private. A recent inequity responsibility account has been posited to explain why people would be impartial in private and suggests that people are fair in private because they want to avoid being responsible for creating partial inequity between others (Gordon-Hecker, Rosensaft-Eshel, et al., 2017; Gordon-Hecker, Choshen-Hillel, Shalvi, & Bereby-Meyer, 2017). Whereas we would agree that people want to avoid being responsible for partiality also in private, it seems that such behavior can be reconciled within the partiality account. The best possible strategy for avoiding appearing partial is to be impartial even in so-called “private” situations (Shaw, 2016). Indeed, there are rarely situations in which it is completely sure that no one could find out one's decision, and even if such situations rarely exist, the allocator herself knows what she did and this could potentially slip out at a later date. People also may initially engage in such behaviors to avoid being judged for being partial and then internalize the norm. Either way, the partiality account would expect decision makers to endorse impartial decision-making even in private circumstances, particularly in cases where the benefits of being partial are low. For example, in resource allocation decisions that occur in third party contexts, the partiality account predicts that people should rarely be partial, even in private, and that there would be very little difference between public and private allocations. The reason is that the benefits of partiality in such cases are low (the allocator may have nothing to gain from giving more to one employee then the other) and the costs of being partial (e.g., condemnation from others) will often outweigh these benefits. Indeed, in such purely third party cases and other cases in which the benefits of partiality are low, people may demonstrate a preference for avoiding partiality (Choshen-Hillel, Shaw, & Caruso,

2015; Gordon-Hecker, Rosensaft-Eshel, et al., 2017; Shaw, 2013). However, when the benefits of partiality increase (e.g., when the potential beneficiary is the allocator or allocator's friend), the partiality account predicts that people will be much more partial and that there should be a larger difference between their public and private allocations. Here people may or may not risk partiality in public, but they may be more likely to be partial in private. Broadly, if the benefits of favoritism increase, we would expect to see more favoritism; if the costs of being judged as partial increase, we would expect to see less favoritism, particularly in public.

10.2. The balance between partiality and favoritism

Our results should not be taken to indicate that people always attempt to avoid partiality. Although people do sometimes feel pressure to be fair, there is a plethora of evidence that people sometimes favor their friends even in cases where this will be seen as unfair (Reinsch & Gardner, 2014; Waytz, Dungan, & Young, 2013). Indeed, people expect such favoritism from their friends (DeScioli & Kurzban, 2011) and will respond negatively to impartial neutrality from their friends (Shaw, Descioli, Barakzai, & Kurzban, 2017). In a supplemental study, we found that allocators were biased in favor of their friends when they made private decisions (see supplemental Study S1). We suspect that in real world contexts, allocators might be even more biased in favor of their friends (especially in private) because they can likely come up with plausible justifications for why their friend deserves the better reward; people are quite adept at justifying their own biased decision-making (DeScioli, Massenkoff, Shaw, Petersen, & Kurzban, 2014; Hsee, 1996; Kunda, 1990; Messick & Sentis, 1979; Shalvi, Gino, Barkan, & Ayal, 2015). There are even formal rules about recusal that stop people from making judgments where they will be biased or worried that they might appear biased (Frost, 2005). Indeed, we found that people opt to recuse themselves when they were concerned with appearing biased or making a biased decision against their friend.

It is worth noting that in some circumstances, partiality may not only be acceptable, but being impartial would be quite strange. Perhaps the clearest example of partiality being acceptable occurs in parent and child relationships: It is certainly acceptable (if not obligatory) to be partial toward one's children in many contexts—a person who felt an equal obligation to feed a stranger as their own child and who correspondingly did not provide their child with everything it needed would not be considered as a very good parent.

Relatedly, even in the office contexts we studied here, there might be cultural and individual differences in whether or not such actions are considered partial or impartial as well as fair or unfair. Although most societies have at least some contexts where they think partiality is unacceptable, especially in the context of legal proceedings (for review, see DeScioli & Kurzban, 2009b), there is considerable variation in what people view as partial and unfair across societies. There is evidence that Western market-based economies place a particularly strong emphasis on valuing fairness and impartiality (Henrich et al., 2010). In many countries, nepotism and cronyism are rampant and may not be considered to be quite as unfair (e.g., Arasli & Tumer, 2008; Brody, Coulter, & Mihalek, 1998; Fiske & Rai, 2014). Such cultural variation is not at odds with the partiality account. From the partiality account perspective, the interesting cross-cultural question is whether people think that behaviors that are unfair are also partial (for a recent investigation on this topic, see Salmon & Serra, 2017). In some of these societies, people may see nepotism ("It is not partial to give more to my brother") as an acceptable justification for inequality in the way that Western people see merit ("It is not partial to give more to the person who did a better job"). Even within the United States, there are differences between what liberals and conservatives consider to be "partial" (Haidt & Graham, 2007); for example, conservatives may see affirmative action as partial because it gives underrepresented minorities an advantage over others and liberals may see not implementing affirmative action as partial because the system is already biased against underrepresented minorities. Furthermore there are likely important individual differences in what, where, and when people think that others should or should not be partial, based on measures like Social value orientation (Van Lange, 1999), empathy (Batson & Ahmad, 2001), and Machiavellianism (Lopes & Fletcher, 2004). Future research will be necessary to investigate the kinds of rules that justify inequality across cultures as well as the individual differences that drive people's judgments within a culture.

10.3. Final remarks

Impartiality is often an antidote for corruption and bias (Dungan, Waytz, & Young, 2014; Shaw & Knobe, 2013). Still, avoiding partiality is no panacea, and the current studies reveal that in some circumstances trying to avoid the appearance of bias can lead people to yet make biased decisions. The current findings are related to similar unanticipated consequences of trying to avoid bias based on ingroup favoritism or conflicts of interests. Recent work demonstrates that when people are told to be "fair" this can actually increase their bias against the dominant ingroup (e.g., white males) rather than actually make them fair (Self, Mitchell, Mellers, Tetlock, & Hildreth, 2015). Another related line of work demonstrates the counterintuitive bias that can occur when someone discloses a conflict of interest. When a person has a conflict of interest that would undermine their objectivity, they (and others) often think they should disclose such conflicts of interest to mitigate this bias (for review see, Cain, Loewenstein, & Moore, 2005). However, such disclosure can perversely lead to more bias because the agent now feels that it is more acceptable to be biased and the other party feels more trusting of the agent who disclosed this conflict (Cain, Loewenstein, & Moore, 2005, 2010; Sah, Loewenstein, & Cain, 2013).

These results do not mean that policy makers should abandon efforts to reduce bias, but that policy makers must understand these consequences and try to design more effective interventions in light of these facts about human psychology. Understanding the ways that our decisions can be biased and debiased is essential for optimizing the distribution of resources in our complex social world.

Open practices

The data for all studies, including supplemental studies, can be found at (https://osf.io/etrcd/). We also have pre-registrations for Study 5 (https://aspredicted.org/zm4d5.pdf) and Study 6 (https://aspredicted.org/gz4bk.pdf).

Acknowledgment

This work was partially supported by a research grant from the Social Enterprise Initiative at The University of Chicago Booth School of Business. Shaw and Choshen-Hillel contributed equally to this work; they knew that each of them might bias against themselves if asked to choose who did a slightly better job, so they ended up choosing the equal option.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jesp.2018.05.009.

References

Abbink, K., Irlenbusch, B., & Renner, E. (2002). An experimental bribery game. *Journal of Law, Economics, and Organization, 18*, 428–454.

Adams, J. (1965). Inequity in social exchange. In L. Berkowitz (Vol. Ed.), *Advances in experimental social psychology. Vol. 2. Advances in experimental social psychology* (pp. 267–299). New York: Academic Press.

Andreoni, J., & Bernheim, B. D. (2009). Social image and the 50–50 norm: A theoretical

and experimental analysis of audience effects. *Econometrica, 77*, 1607–1636.
Arasli, H., & Tumer, M. (2008). Nepotism, favoritism and cronyism: A study of their effects on job stress and job satisfaction in the banking industry of north Cyprus. *Social Behavior and Personality: An International Journal, 36*, 1237–1250.
Balliet, D., Wu, J., & De Dreu, C. K. W. (2014). Ingroup favoritism in cooperation: A meta-analysis. *Psychological Bulletin, 140*, 1556–1581.
Barakzai, A., & Shaw, A. (2018). Friends without benefits: When "helpful" friends are worse than unhelpful friends. *Evolution and Human Behavior* (in press).
Barr, A., & Serra, D. (2009). The effects of externalities and framing on bribery in a petty corruption experiment. *Experimental Economics, 12*, 488–503.
Batson, C. D., & Ahmad, N. (2001). Empathy-induced altruism in a prisoner's dilemma II: What if the target of empathy has defected? *European Journal of Social Psychology, 31*, 25–36.
Batson, C. D., Kobrynowicz, D., Dinnerstein, J. L., Kampf, H. C., & Wilson, A. D. (1997). In a very different voice: Unmasking moral hypocrisy. *Journal of Personality and Social Psychology, 72*, 1335–1348.
Brewer, M. B., & Silver, M. (1978). Ingroup bias as a function of task characteristics. *European Journal of Social Psychology, 8*, 393–400.
Brick, I., Palmon, O., & Wald, J. (2006). CEO compensation, director compensation and firm performance: Evidence of cronyism? *Journal of Corporate Finance, 12*, 403–423.
Brody, R. G., Coulter, J. M., & Mihalek, P. H. (1998). Whistle-blowing: A cross-cultural comparison of ethical perceptions of U.S. and Japanese accounting students. *American Business Review, 16*, 14–21.
Cain, D. M., Loewenstein, G., & Moore, D. A. (2005). The dirt on coming clean: Perverse effects of disclosing conflicts of interest. *The Journal of Legal Studies, 34*, 1–25.
Cain, D. M., Loewenstein, G., & Moore, D. A. (2010). When sunlight fails to disinfect: Understanding the perverse effects of disclosing conflicts of interest. *Journal of Consumer Research, 37*, 836–857.
Choshen-Hillel, S., Shaw, A., & Caruso, E. M. (2015). Waste management: How reducing partiality can promote efficient resource allocation. *Journal of Personality and Social Psychology, 119*, 210–231.
Choshen-Hillel, S., Shaw, A., & Caruso, E. M. (2018). Disadvantaged but not dissatisfied: Agency ameliorates negative reactions to unequal pay. *Journal of Experimental Psychology: Applied* (in press).
Choshen-Hillel, S., & Yaniv, I. (2011). Agency and the construction of social preference: Between inequality aversion and prosocial behavior. *Journal of Personality and Social Psychology, 101*, 1253–1261.
Clark, M. S., & Mills, J. (1979). Interpersonal attraction in exchange and communal relationships. *Journal of Personality and Social Psychology, 37*, 12–24.
Cole, T., & Teboul, J. C. B. (2004). Non-zero-sum collaboration, reciprocity, and the preference for similarity: Developing and adaptive model of close relational functioning. *Personal Relationships, 11*, 135–160.
Cooney, G., Gilbert, D. T., & Wilson, T. D. (2016). When fairness matters less than we expect. *Proceedings of the National Academy of Sciences, 113*, 11168–11171.
Cox, J. (2004). How to identify trust and reciprocity. *Games and Economic Behavior, 46*, 260–281.
Damon, W. (1977). *The social world of the child.* San Francisco: Jossey-Bass.
Dana, J., Cain, D. M., & Dawes, R. M. (2006). What you don't know won't hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes, 100*, 193–201.
Dawes, C., Fowler, J., Johnson, T., McElreath, R., & Smirnov, O. (2007). Egalitarian motives in humans. *Nature, 446*, 794–796.
Delton, A. W., Krasnow, M. M., Cosmides, L., & Tooby, J. (2011). Evolution of direct reciprocity under uncertainty can explain human generosity in one-shot encounter. *Proceedings of the National Academy of Sciences, 108*, 13335–13340.
DeScioli, P., & Kurzban, R. (2009a). The alliance hypothesis for human friendship. *PLoS One, 4*, e5802.
DeScioli, P., & Kurzban, R. (2009b). Mysteries of morality. *Cognition, 112*, 281–299.
DeScioli, P., & Kurzban, R. (2011). The company you keep: Friendship decisions from a functional perspective. In J. I. Krueger (Ed.). *Social judgment and decision making* (pp. 205–224). New York: Psychology Press.
DeScioli, P., & Kurzban, R. (2013). A solution to the mysteries of morality. *Psychological Bulletin, 139*, 477–496.
DeScioli, P., Kurzban, R., Koch, E. N., & Liben-Nowell, D. (2011). Best friends, alliances, friend ranking, and the MySpace social network. *Perspectives on Psychological Science, 6*, 6–8.
DeScioli, P., Massenkoff, M., Shaw, A., Petersen, M. B., & Kurzban, R. (2014). Equity or equality? Moral judgments follow the money. *Proceedings of the Royal Society of London B: Biological Sciences, 281*, 2014–2112.
Dungan, J., Waytz, A., & Young, L. (2014). Corruption in the context of moral trade-offs. *Journal of Interdisciplinary Economics, 26*, 97–118.
Fiske, A. P. (1992). The four elementary forms of sociality: Framework for a unified theory of social relations. *Psychological Review, 99*, 689–723.
Fiske, A. P., & Rai, T. S. (2014). *Virtuous violence: Hurting and killing to create, sustain, end, and honor social relationships.* Cambridge University Press.
Frost, A. (2005). Keeping up appearances: A process-oriented approach to judicial recusal. *University of Kanas Law Review, 53*, 531–593.
Gordon-Hecker, T., Choshen-Hillel, S., Shalvi, S., & Bereby-Meyer, Y. (2017). Resource allocation decisions: When do we sacrifice efficiency in the name of equity? In M. Li, & D. Tracer (Eds.). *Interdisciplinary perspectives on fairness, equity and justice* (pp. 93–105). NY: Springer.
Gordon-Hecker, T., Rosensaft-Eshel, D., Pittarello, A., Shalvi, S., & Bereby-Meyer, Y. (2017). Not taking responsibility: Equity trumps efficiency in allocation decisions. *Journal of Experimental Psychology: General, 146*, 771–775.
Gurven, M. (2004). To give and to give not: The behavioral ecology of human food transfers. *Behavioral and Brain Sciences, 27*, 543–559.
Gurven, M. (2006). The evolution of contingent cooperation. *Current Anthropology, 47*, 185–192.
Haidt, J., & Graham, J. (2007). When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research, 20*, 98–116.
Henrich, J., Ensminger, J., Mcelreath, R., Barr, A., Barrett, C., Bolyanatz, A., ... Lesorogol, C. (2010). Markets, religion, community size, and the evolution of fairness and punishment. *Science, 327*, 1480–1484.
Hook, J., & Cook, T. D. (1979). Equity theory and the cognitive ability of children. *Psychological Bulletin, 86*, 429–445.
Hsee, C. K. (1996). Elastic justification: How unjustifiable factors influence judgments. *Organizational Behavior and Human Decision Processes, 66*, 122–129.
Hughes, J. S. (2017). In a moral dilemma, choose the one you love: Impartial actors are seen as less moral than partial ones. *British Journal of Social Psychology, 56*, 561–577.
Kenny, D. A., Mohr, C. D., & Levesque, M. J. (2001). A social relations variance partitioning of dyadic behavior. *Psychological Bulletin, 127*, 128–141.
Kleiman-Weiner, M., Shaw, A., & Tenenbaum, J. B. (2017). Constructing social preferences from anticipated judgments: When impartial inequity is fair and why? *Proceedings of the 39th annual conference of the cognitive science society*.
Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin, 108*, 480–498.
Laurenceau, J. P., Barrett, L. F., & Pietromonaco, P. R. (1998). Intimacy as an interpersonal process: The importance of self-disclosure, partner disclosure, and perceived partner responsiveness in interpersonal exchanges. *Journal of Personality and Social Psychology, 74*, 1238–1251.
Leary, M. R., Allen, A. B., & Terry, M. L. (2011). Managing social images in naturalistic versus laboratory settings: Implications for understanding and studying self-presentation. *European Journal of Social Psychology, 41*, 411–421.
Levitt, S., & List, J. (2007). What do laboratory experiments measuring social preferences tell us about the real world? *The Journal of Economic Perspectives, 21*, 153–174.
Liberman, Z., & Shaw, A. (2017). Children use partial sharing as a cue to friendship. *Journal of Experimental Child Psychology, 159*, 97–109.
Lopes, J., & Fletcher, C. (2004). Fairness of impression management in employment interviews: A cross-country study of the role of equity and Machiavellianism. *Social Behavior and Personality: An International Journal, 32*, 747–768.
MacKinnon, D. P., Fairchild, A. J., & Fritz, M. S. (2007). Mediation analysis. *Annual Review of Psychology, 58*, 593–614.
Messick, D. M. (1995). Equality, fairness, and social conflict. *Social Justice Research, 8*, 153–173.
Messick, D. M., & Sentis, K. P. (1979). Fairness and preference. *Journal of Experimental Social Psychology, 15*, 418–434.
Mitchell, G., Tetlock, P. E., Mellers, B. A., & Ordonez, L. D. (1993). Judgments of social justice: Compromises between equality and efficiency. *Journal of Personality and Social Psychology, 65*, 629–639.
Niemi, L., Wasserman, E., & Young, L. (2018). The behavioral and neural signatures of distinct conceptions of fairness. *Social Neuroscience, 13*, 399–415.
Piazza, J., & Bering, J. M. (2008). Concerns about reputation via gossip promote generous allocations in an economic game. *Evolution and Human Behavior, 29*, 172–178.
Pietraszewski, D. (2016). How the mind sees coalitional and group conflict: The evolutionary invariances of n-person conflict dynamics. *Evolution and Human Behavior, 37*, 470–480.
Rand, D. G., Arbesman, S., & Christakis, N. A. (2011). Dynamic networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences, 108*, 19193–19198.
Reinsch, N. L., & Gardner, J. A. (2014). Do communication abilities affect promotion decisions? Some data from the C-Suite. *Journal of Business and Technical Communication, 28*, 31–57.
Reis, H. T., & Gruzen, J. (1976). On mediating equity, equality, and self-interest: The role of self-presentation in social exchange. *Journal of Experimental Social Psychology, 12*, 487–503.
Sah, S., Loewenstein, G., & Cain, D. M. (2013). The burden of disclosure: Increased compliance with distrusted advice. *Journal of Personality and Social Psychology, 104*, 289–304.
Salmon, T. C., & Serra, D. (2017). Corruption, social judgment and culture: An experiment. *Journal of Economic Behavior and Organization, 142*, 64–78.
Self, W. T., Mitchell, G., Mellers, B. A., Tetlock, P. E., & Hildreth, J. A. D. (2015). Balancing fairness and efficiency: The impact of identity-blind and identity-conscious accountability on applicant screening. *PLoS One, 10*, e0145208.
Shalvi, S., Gino, F., Barkan, R., & Ayal, S. (2015). Self-serving justifications: Doing wrong and feeling moral. *Current Directions in Psychological Science, 24*, 125–130.
Shaw, A. (2013). Beyond "To share or not to share": The impartiality account of fairness. *Current Directions in Psychological Science, 22*, 413–417.
Shaw, A. (2016). Fairness: What it isn't, what it is, and what it might be for. *Evolutionary perspectives on child development and education* (pp. 193–214). Springer International Publishing.
Shaw, A., & Choshen-Hillel, S. (2017). It's not fair: Folk intuitions about disadvantageous and advantageous inequity aversion. *Judgment and Decision making, 12*, 208–223.
Shaw, A., Choshen-Hillel, S., & Caruso, E. (2016). The development of inequity aversion: Understanding when (and why) people give others the bigger piece of the pie. *Psychological Science, 27*, 1352–1359.
Shaw, A., Descioli, P., Barakzai, A., & Kurzban, K. (2017). Whoever is not with me is against me: The costs of neutrality among friends. *Journal of Experimental Social Psychology, 71*, 96–104.
Shaw, A., Descioli, P., & Olson, K. R. (2012). Fairness versus favoritism in children. *Evolution and Human Behavior, 33*, 736–745.
Shaw, A., & Knobe, J. (2013). Not all mutualism is fair and not all fairness is mutualistic. *Behavioral and Brain Sciences, 36*, 100–101.
Shaw, A., Montinari, N., Piovesan, M., Olson, K. R., Gino, F., & Norton, M. I. (2014).

Children develop a veil of fairness. *Journal of Experimental Psychology: General, 143*, 363–375.
Shaw, A., & Olson, K. R. (2012). Children discard a resource to avoid inequity. *Journal of Experimental Psychology: General, 141*, 382–395.
Shaw, A., & Olson, K. R. (2014). Fairness as an aversion to partiality: The development of procedural justice. *Journal of Experimental Child Psychology, 119*, 40–53.
Sweeney, P. D., & McFarlin, D. B. (2004). Social comparisons and income satisfaction: A cross- national examination. *Journal of Occupational and Organizational Psychology, 77*, 149–154.
Tajfel, H., Billig, M. G., Bundy, R. P., & Flament, C. (1971). Social categorization and intergroup behaviour. *European Journal of Social Psychology, 1*, 149–178.
Thibaut, J. W., & Walker, L. (1975). *Procedural justice: A psychological analysis.* (L. Erlbaum Associates).
Tooby, J., & Cosmides, L. (1996). Friendship and the Banker's Paradox: Other pathways to the evolution of adaptations for altruism. In W. G. Runciman, J. Maynard Smith, & R. I. M. Dunbar (Vol. Eds.), *Evolution of social behaviour patterns in primates and man. Proceedings of the British Academy. Vol. 88. Evolution of social behaviour patterns in primates and man. Proceedings of the British Academy* (pp. 119–143).
Trivers, R. (1971). Evolution of reciprocal altruism. *Quarterly Review of Biology, 46*, 35–57.
Tyler, T. R. (2000). Social justice. *International Journal of Psychology, 35*, 117–125.
Van Lange, P. A. M. (1999). The pursuit of joint outcomes and equality in outcomes: An integrative model of social value orientation. *Journal of Personality and Social Psychology, 77*, 337–349.
Waytz, A., Dungan, J., & Young, L. (2013). The whistleblower's dilemma and the fairness-loyalty tradeoff. *Journal of Experimental Social Psychology, 49*, 1027–1033.