On the Merits of Separate Spaces: Why Institutions Isolate Cooperation and Division Tasks

This version: May 24, 2023

Scott E. Bokemper[1] and Gregory A. Huber[2,3,4]*

[1] Independent Researcher, Waukee, Iowa, United States of America

[2] Institution for Social and Policy Studies, Yale University, New Haven, Connecticut, United States of America

[3] Center for the Study of American Politics, Yale University, New Haven, Connecticut, United States of America

[4] Department of Political Science, Yale University, New Haven, Connecticut, United States of American

* Corresponding author

E-mail: gregory.huber@yale.edu

ISPS/Yale University, 77 Prospect Street, New Haven CT 06511

Funding: Not applicable

Conflict of interest: The authors declare no conflicts of interest regarding this work.

Availability of data and replication materials: Data and code necessary to reproduce the analysis presented in the manuscript is available at the Harvard Dataverse https://doi.org/10.7910/DVN/H7BWDJ.

Ethics Approval: The studies presented here were approved by the Yale University IRB.

Consent to Participate and Publication: Written consent was obtained prior to participation. The consent form informed participants that results may be used in publication.

Acknowledgments: We thank the Center for the Study of American Politics and the Institution for Social and Policy Studies at Yale for financial support. Additionally, we thank Chris Dawes and participants in the American Politics Workshop at Yale University for helpful comments and advice.

On the Merits of Separate Spaces: Why Institutions Isolate Cooperation and Division Tasks

Abstract
Do institutions shape the possibility of sustaining cooperation when the same individuals must first divide resources and then attempt to cooperate? It could be that simply having received an inequitable division undermines cooperative behavior, reducing aggregate welfare. Alternatively, it might be that only when interacting with the same individual or group does this spillover occur, in which case separating tasks across institutions may prevent this negative spillover. To test these arguments, we designed a two-stage incentivized experiment in which participants interact in a division task and then in a task in which cooperation improves aggregate welfare. In two experiments, individuals were randomly assigned to interact either with the same individual for both tasks or with a different individual for each task. In the second experiment, individuals could also interact with a person who was in the same arbitrary group as their partner in the division task. Holding constant both past history and past partner behavior, the results of these experiments provide support for a Partner History effect in which the mechanism that produces spillover is interacting with the same individual in both decisions. We also find evidence for a weaker Group History effect in which negative spillover occurs when the partner in the cooperative task is a member of the same group as the partner from the division task.

Word count: 8950
Keywords: behavioral spillover, cooperation, experiment

Political and economic institutions can facilitate cooperation by structuring incentives to suppress selfish behavior. For example, institutions that sanction free riders make the practice less rewarding, bolstering cooperation (e.g. Chang, Dawes, & Johnson, 2018). Other institutions provide information about past behavior that signals whether someone can be trusted to reciprocate cooperation (e.g. Milgrom, North, & Weingast, 1990). Cooperative interactions are positive-sum with at least one party being made better off. In contrast, conflictual games, like the division of resources, pit individuals against each other with one gaining at the other's expense. (And, of course, many interactions are mixed-motive.)

Many political actors, such as nation states, members of Congress, hierarchies (e.g., administrative agencies), and local governments, are responsible for making both conflictual and cooperative decisions, as are small groups. Focusing on elite interactions, a city council might first have to make a rival decision about locating a valuable resource like a park or an eyesore like a dump before turning to a cooperative decision like raising taxes for citywide road repair or to improve public schools. Does the outcome of the rival decision affect the council's subsequent ability to cooperate? What if the tasks can be separated, for example if a committee that makes the location decision can be separated from the committee that proposes the citywide taxing policy? An important theoretical question is therefore whether and why "losing" in the division of resources causes an individual to be less cooperative when there is an opportunity for a group of individuals to benefit in a cooperative interaction. If spillovers arise, can institutions reduce these potential threats to cooperation?

A burgeoning literature in the behavioral sciences has examined when spillover occurs in a variety of contexts. Yet, in many instances it is unclear what the exact mechanism is that causes spillover across games. Is it the outcome that an individual experienced in the first context, their payoff history? Is it their partner's reputation? Or is it who the individual was interacting with across contexts? We propose that a mechanism that drives spillover from division tasks to cooperation tasks is a Partner History effect. Specifically, experiencing a loss in a division task will reduce subsequent cooperation if both tasks take place with the same individual. This effect is distinct from the effect of experiencing a loss (one's own history) and the effect of interacting with an individual who was uncooperative in the past (the partner's reputation). In other words, there is something unique about sharing a personal history with an individual that cannot be replicated even when all relevant information and experience are present.

This theoretical account focuses on distinguishing the effect of a prior partner from a new partner, ignoring the fact that many elite and mass interactions are also structured by group identities such as partisanship, race or ethnicity, or other factors (Sheffer and Loewen 2019). This leads naturally to the question of whether spillover also arises when choosing to cooperate with someone who is a member of the same group as a prior partner in a division task, such as when a legislator has to choose whether to cooperate with an out-partisan whose fellow out-partisan was able to influence the legislator's experience in a prior division of resources. In short, is there also a Group History effect? Understanding this is important because it has implications for the design of institutions that seek to reduce the possibility of spillover. For example, a legislative committee system that separates rival and cooperative tasks may be ineffective in precluding spillovers if people carry forward their negative reaction to an out-partisan's (selfish) behavior in interactions with other out-partisans.

To examine behavioral spillover across rival and cooperative tasks empirically, we designed a two-stage, incentivized game. In the first stage, participants played the Division Game. This was a conflictual game with zero-sum payoffs: For every token that one player gained, the other player lost. In the second stage, participants played the Policy Game, which was a cooperative (positive-sum) game in which one player could enhance the welfare of the other and, in some instances, also benefit themselves. Using this framework, we conducted two experiments to study the conditions under which receiving an unequal (unfair or inequitable) allocation in the Division Game, versus an equal (or fair) allocation, reduced support for adopting a Pareto-improving policy in the Policy Game. In the second experiment, we added conditions that allow us to assess whether spillover would also occur in the presence of shared group membership.

Background and Motivation

Institutions that combine division tasks and cooperative tasks cause the same individuals to interact repeatedly. This type of interaction makes direct reciprocity possible. Interacting with the same person repeatedly allows individuals to establish reputations for being cooperative and as such all parties are better off as long as they avoid the short-run temptation to defect (Axelrod & Hamilton, 1981). When defection does occur, strategies like tit-for-tat call for defection by a partner to be met with defection in the next period. Reciprocity can also occur in sequential games. For instance, in two-player ultimatum games, receivers commonly reject offers that are less than 30% of the total endowment leaving both players with nothing (Camerer & Thaler, 1995). Similarly, participants in a one-shot public goods game engaged in costly punishment of those who contributed less than they did to the public good even though punishment reduced efficiency and many participants who were punished had made Pareto-improving contributions to the public good (Bicskei, Lankau, & Bizer, 2016).

Past work has examined how interacting with the same individual(s) sequentially across contexts can affect cooperation.[1] For instance, participants who had a history of successful coordination together were more likely to cooperate in a subsequent prisoner's dilemma than participants who did not (Knez & Camerer, 2000). Other research has found that behavior spills over from one coordination game to another and from one average opinion game to another (e.g., Cason, Savikhin, & Sheremeta, 2012). An explanation for these findings is that behavior in the first interaction establishes a precedent for cooperative behavior that signals a willingness to cooperate in a different context with distinct strategic incentives (Knez & Camerer, 2000).

But repeatedly interacting with the same individual confounds three things: 1) the effect of a partner's past behavior, e.g. their reputation, 2) the effect of experiencing an outcome (that is, a person's experiential history), and 3) the effect of a shared history with the specific person who caused that outcome. It is that last theoretical effect that we wish to isolate empirically. Disentangling these effects has potentially important consequences for the design of institutions that seek to sustain cooperation. On one hand, if the effect of experiencing a bad outcome predominates, institutions that separate responsibility for division and cooperation will not mitigate spillover because an individual will still experience a loss while others will experience a win in the division of resources. Similarly, if it is a partner's past behavior that matters because it establishes their reputation, who that partner was previously interacting with is irrelevant. On the other hand, if the mechanism driving spillover is repeated interaction with the *same individual*, dividing responsibility by creating different decision partners in each context may prevent conflict over the division of resources from undermining future cooperation. That is, institutions that isolate division and cooperation tasks may enhance future cooperation.

[1] Although our focus here is sequential interaction, past research on spillover has also addressed simultaneous interaction in two different strategic contexts (e.g. Bednar, Chen, Liu, & Page, 2012).

Indirect Reciprocity, Experiencing an Outcome, and Reputation

Institutions that separate division tasks from cooperation tasks break apart a shared history of play between parties; however, this does not erase the outcome that was experienced nor does it eliminate a new party's past behavior toward an unrelated party. If spillover occurs in these situations, it is potentially driven by indirect reciprocity in which Individual A takes a selfish (cooperative) action towards Individual B because Individual B took a selfish (cooperative) action towards individual C.[2] That is, such indirect reciprocity may arise because of a partner's reputation. Past research has found that the ability to establish and track reputation is necessary for indirect reciprocity to occur. For instance, participants were more likely to help another participant in a helping game when information about their prior helpfulness in the previous five periods was made public instead of kept private (Engelmann & Fischbacher, 2009).

Extant research examining spillover across contexts when interacting with different individuals has found that initial cooperative behavior begets future cooperative behavior. For instance, Ahn et al. (2001) found an association between first-stage coordination and second-stage cooperation in a prisoner's dilemma when participants were randomly matched with a new individual in the second stage, but this spillover effect is weaker than when interacting with the same individual in both stages. Other work has found that people who were more generous to another individual received larger transfers in a subsequent trust game (Jordan et. al, 2016). Further, people who gave more to a charity were cooperated with more frequently in both a prisoner's dilemma and a trust game (Albert et al., 2007). Thus, knowledge of how individuals have behaved in the past can promote future cooperation, even across different strategic contexts.

[2] This could also be due a different kind of reciprocity, upstream (generalized) reciprocity (Nowak & Roch, 2007). Our design cannot disentangle these different motivations, but instead holds constant past outcomes to allow us to isolate the relative importance of that outcome originating with the same vs. different partners.

But other experiments have failed to produce spillover from one context to another when interacting with different individuals. Capraro and Marcelletti (2014) did not find that people cooperated more with a stranger in a prisoner's dilemma game or a dictator game after having received their endowment from the generous action of a different individual compared to those who received their endowment from the experimenter. This implies that an individual's personal history does not alter how they behave towards a new partner in subsequent interactions.

Cumulatively, this past experimental work does not directly assess the causal effect of interacting with different individuals compared to the same individual on spillover across strategic contexts. That is, spillover may not occur when someone interacts with different individuals but may occur in situations in which individuals interact repeatedly together. Here, we hold constant the experience that an individual has and their partner's past behavior to directly test how varying pairing protocols affects (or creates) spillover across contexts. In contrast to previous work on indirect reciprocity, aligning a new partner's past behavior with an old partner's past behavior allows for an unconfounded test of how partner matching affects behavioral spillover.

Group History and Spillover

A potentially relevant, but currently unexplored, factor that could affect spillover is the role of group affiliation. One possibility is that the hypothesized Partner History effect operates only through interactions with that person. But many interactions involve individuals who are members of larger groups. Extant research has shown that group membership, even if it is established arbitrarily, can have substantial effects on social behavior (e.g. Balliet, Wu, & De Dreu, 2014). Unfair treatment of individuals on the basis of group membership is perceived as less fair than random mistreatment (Bokemper, DeScioli, and Kline 2019), and group membership also affects how leaders are held accountable (Landa and Duell 2015). Further, past experiments have found preferential treatment of in-group members compared to out-group members and unaffiliated individuals in a plethora of economic experiments (e.g. Fowler & Kam, 2007).

How does an individual respond to a loss caused by an out-group member? More specifically, will an individual who experienced a loss at the hands of one out-group member take a costly action against a different out-group member? When interacting with the exact same individual, a person is given the opportunity to retaliate against (reward) that specific person for an undesirable (desirable) outcome, a reciprocity account that undergirds the predicted Personal History effect. Again, if the experience is what drives spillover between division and cooperation tasks, it is reasonable to expect that receiving an unfair division from one group member will reduce their willingness to take an action that benefits a different out-group member. However, while shared group membership, even when created artificially, increases cooperative behavior, such ingroup bias is reduced when there is explicit evidence about someone's willingness to cooperate (e.g., Misch, Pauls, and Dunham 2021). Consequently, it is unclear whether spillover is likely to occur on the basis of a prior partner's shared group membership with a current partner when information about the new partner's behavior is available and held constant. In other words, if sharing a history with an individual is the key mechanism, then two interactions with different out-group members will be treated as separable and spillover will not be observed. Extant research does not provide clear expectations about how much of a personal history effect transfers to a different group member, a Group History effect.

Expectations

In the experiments presented below, we test three core predictions focusing on when losing in the division of resources will cause reduced subsequent cooperation. Drawing on the above arguments, we identify three specific expectations. First:

General Experience Effect: Receiving an unfair (inequitable) division rather than a fair (equitable) division decreases support for a cooperative policy that benefits someone who made an unfair allocation to another person.

We view this effect as encompassing: an individual's endowment, an individual's motivation to be selfish towards someone who was selfish towards a third party (indirect reciprocity), and an individual's willingness to pay forward a selfish action that was inflicted on them (displaced reciprocity). Neither of these effects is our core theoretical focus, but as we note above, both are conflated empirically with the Personal History effect:

Personal History Effect: Receiving an unfair division from the same person decreases support for a cooperative policy that benefits that other person.

In this account, an individual will not cooperate with someone who was unfair to them compared to someone who was unfair to another person (as in the General Experience Effect). Notably, this effect relies on that specific person having behaved in a particular way toward me in prior interactions. This effect is in addition to the General Experience Effect that would also exist in any repeated interaction with the same person.

Finally, we examine the:

Group History Effect: Receiving an unfair division from a member of an out-group decreases support for a cooperative policy that benefits a different member of that out-group compared to when those individuals are members of different groups.

In this account, the selfish actions of one out-group member caused an individual to take an action that negatively affects a different out-group member. Once again, this effect is in addition to the General Experience Effect.

Experimental Design

We test these expectations with a two-stage experiment that is a division task followed by a cooperative task. Unless otherwise noted, participants were fully informed about all features of the games described below. Participants also knew that their decisions had monetary consequences for themselves and others.

Stage 1: The Division Game

In the Stage 1 Division Game, participants were paired with another player and each proposed a division of a common pot of tokens. This game is conflictual (zero-sum) because each gain for one player caused a loss for the other player. This task proceeded in steps and players were told about all steps before the task began. First, each player received an initial endowment of 10 tokens. Then, both players put their tokens into a common pot that was doubled by the experimenter, resulting in a total pot of 40 tokens.

Next, each player made a private proposal to divide the pot of tokens from a menu of specific options. Players could propose to take all 40 tokens (leave none for their partner), take 30 tokens (leave 10 for their partner, only in Experiment 1), or take 20 tokens (leave 20 for their partner). Note that one inequitable division would leave the partner no worse off than their initial allocation of 10 tokens. After both players made a proposal, the computer deterministically chose one player to be the "Decider" and their private proposal was made public and implemented resulting in an outcome for both the Decider and the Receiver. Consequently, players made their proposal before they knew whether it would be chosen and only learned about their partner's proposal after they had made their own.

At the conclusion of the game (after outcomes were realized) participants proceeded to the next stage of the experiment, the Policy Game. Subjects were only informed about this game after the completion of the Division Game. This was done so behavior in the Division Game would not be affected by the expectation of future behavior in a subsequent game, a design feature we discuss in detail in the conclusion.

Stage 2: The Policy Game

In the Policy Game, each player indicated whether they would support a policy that would generate and then divide a common pool of tokens. This task proceeded in stages and players were told about all steps before the task began. First, each player was given a new endowment of 10 tokens. Second, players were informed about the prior behavior and earnings of their partner in the Division Game (see below). Third, players proposed to adopt a policy or simply retain their endowment and one of their proposals was randomly chosen to be implemented. If the policy was adopted, both players paid 10 tokens from their endowment and received a payoff. Participants were randomly assigned to one of three policy conditions shown in Table 1. Note that all policies are (weakly) Pareto improving. Because we fix the payoffs to Player B, we label the policies according to the payoff that Player A would receive if the policy was adopted.

Table 1: Policy Game Net Payoffs from Adopting Policy Compared to Retaining 10 token Endowment by Policy.

<table>
<tr><td>Policy:</td><td>Net 0</td><td>Net 10</td><td>Net 20</td></tr>
<tr><td>Player A</td><td>0 tokens</td><td>10 tokens</td><td>20 tokens</td></tr>
<tr><td>Player B</td><td>20 tokens</td><td>20 tokens</td><td>20 tokens</td></tr>
<tr><td colspan="4">Note: These are Net payoffs relative to simply keeping 10 allocated tokens. So, for example, in the Net 0 condition, Player A contributes 10 tokens and receives 10 tokens back, while Player B contributes 10 tokens and receives 30 tokens back.</td></tr>
</table>

Notably, Player A is never made worse off by adopting any policy, but they derive less of a benefit from the policy than Player B in the Net 0 and Net 10 conditions. For those in the Player A role, we view the Net 0 condition as a particularly good test case for our predictions relating to spillover as many interactions require the cooperation of individuals whose benefit will not exceed their cost while others might benefit greatly. By contrast, when the policy generates a large benefit to an individual (i.e. the Net 20 condition) choosing to not support the policy as a way to punish someone else is very costly, and therefore unlikely. This design is therefore analogous to other research that varies the cost of punishment (e.g., Burnham 2015), and we expect spillover to be less likely in the Net 20 condition.

Before deciding whether to propose adopting the policy or not, players learned about who their partner was in the Policy Game and their partner's behavior and earnings in the Division Game. After learning this information, players made their decisions and which player's decision would be implemented was randomly determined. Thus, every player should cast a vote as if they were pivotal.

We used this framework to conduct two experiments (summarized in Table 2). In Experiment 1, we tested the Partner History effect compared to the General Experience effect. To isolate these effects from each other, we compare a Same Person condition where participants interacted with the same individual for both games to a Duplicate Behavior Stranger (Duplicate hereafter) condition in which participants interacted with two different individuals for each game. Crucially, in the Duplicate condition, a participant's Policy Game partner behaved identically in the Division Game to how their actual Division Game partner behaved, with the only difference being that they behaved that way toward someone else. The General Experience effect is therefore held constant across conditions. In Experiment 2, we tested the Group History effect, the degree to which having a negative experience with one group member spills over and affects behavior towards another. We did this by having the Duplicate condition partners be a member of an unrelated group (baseline Duplicate condition) or the same group (Same Group condition) as the player's Division Task partner.

Table 2: Summary of Procedure and Experimental Design

	Experiment 1	Experiment 2
Platform	Amazon Mechanical Turk	Amazon Mechanical Turk
Incentives	50 cents + bonus payment	50 cents + bonus payment
Dates	November 2018	February/March 2019
Size	N = 987	N = 1,776
Manipulations		
Decision Game Outcome	0 tokens, 10 tokens, or 20 tokens	0 tokens or 20 tokens
Pairing	*Same Person*: play both games with same person, or *Duplicate Person:* play each game with a different individual	*Same Person*: play both games with same out-group member, or *Duplicate Person:* play Division Game with member of Group B and Policy Game with member of Group C, or *Same Group:* play each game with a different member of the same Group B
Policy Game Payoff	Net 0 tokens, Net 10 tokens, or Net 20 tokens	Net 0 tokens, Net 10 tokens, or Net 20 tokens

Experiment 1: Personal History vs. General Experience

We recruited participants from Amazon Mechanical Turk to complete a decision-making study that was administered using the Qualtrics survey platform. Our recruitment filtered for MTurk workers who had at least a 98% lifetime HIT approval rating and were located in the United States. Participants received 50 cents for completing the study and a bonus payment based on the tokens that they earned in the study (1 token = 1 cent). Participants were not told anything about the Policy Game (including that they would play a second game) until after they had completed the Division Game. After reading the instructions for each game, participants took a short comprehension quiz to confirm that they understood what they read. Participants received 2 cents for each comprehension question that they answered correctly to ensure that they read and understood the instructions. 91% of the participants answered at least six of the seven comprehension questions correctly.

Our recruitment procedure was completed in two steps. First, we recruited participants who took the role of Decider in the Division Game and then proposed whether to adopt each of the three policies (or keep their endowments) in the Policy Game. We do not study the behavior of these participants, because we are interested in the effect of the Division Game outcome induced by these players on the Policy Game outcome of other subjects. Instead, these participants are used as partners for pairings in both the Division Game and Policy Game. As with the participants recruited to be Receivers in the Division Game (i.e., not Deciders), these players were compensated based on their choices and the outcome of the Policy Game.[3]

We then recruited a second wave of participants (n = 987) who were Receivers in the Division Game. These participants then took the role of Player A in the Policy Game, meaning they had variation in the cost of choosing not to cooperate with a partner (Player B) whose net payoffs were held constant across policy conditions. These participants', whose behavior is the focus of our study, were randomly assigned to a Division Game allocation, which corresponded to real decisions made by a human player recruited earlier, and a Policy Game environment with a policy (Net 0, Net 10, or Net 20) and a pairing condition: *Same Person:* play both stages with the same person or *Duplicate Person*: play the Division Game and the Policy Game with different persons. Note that regardless of pairing condition, partners in the Policy Game had

[3] Those who were recruited to be the Decider were matched to multiple participants who were Receivers. Because the behavior of the Deciders in the Policy Game is not of theoretical interest, this decision allowed us to implement the experiment for substantially reduced costs.

played the Division Game and made the exact same allocation as a participant's partner did in the Division Game (see Supplementary Table A1 for sample size by condition).

The difference between the Duplicate Person and Same Person conditions allows us to estimate the difference between the General Experience Effect and the Partner History Effect (see Figure 1). Given that participants interacted with a different individual in the Division Game and the Policy Game in the Duplicate Person condition, the effect of the Division Game outcome on the Policy Game could not be caused by how a participant was treated by their *specific* partner in the Division Game. Instead, the effect of the Division Game outcome on the Policy Game would need to originate from the participant's experience in the Division Game or the Policy Game partner's behavior in the Division Game (their reputation) (Figure 1 Panel A). By contrast, in the Same Person condition, the spillover between contexts can arise from how a participant was treated by the specific individual who they are interacting with in both games (Figure 1 Panel B) as well as those other mechanisms. The difference between the Same Person and Duplicate Person conditions therefore can only be attributed to a participant's experience with a specific individual because past experience and partner reputation are held constant across conditions. That is, comparing outcomes across these two conditions allows us to "difference out" causal effects, providing an uncontaminated estimate of the Personal History Effect.

Figure 1: Causal models for effects in Experiment 1
Panel A: General Experience Effect (Duplicate Person condition)

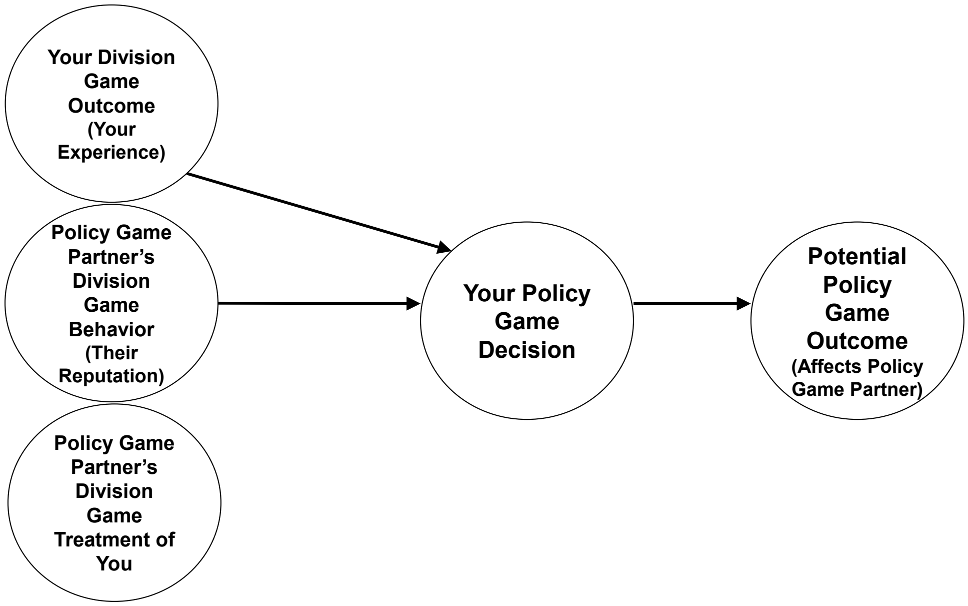

Panel B: Partner History Effect (Same Person condition)

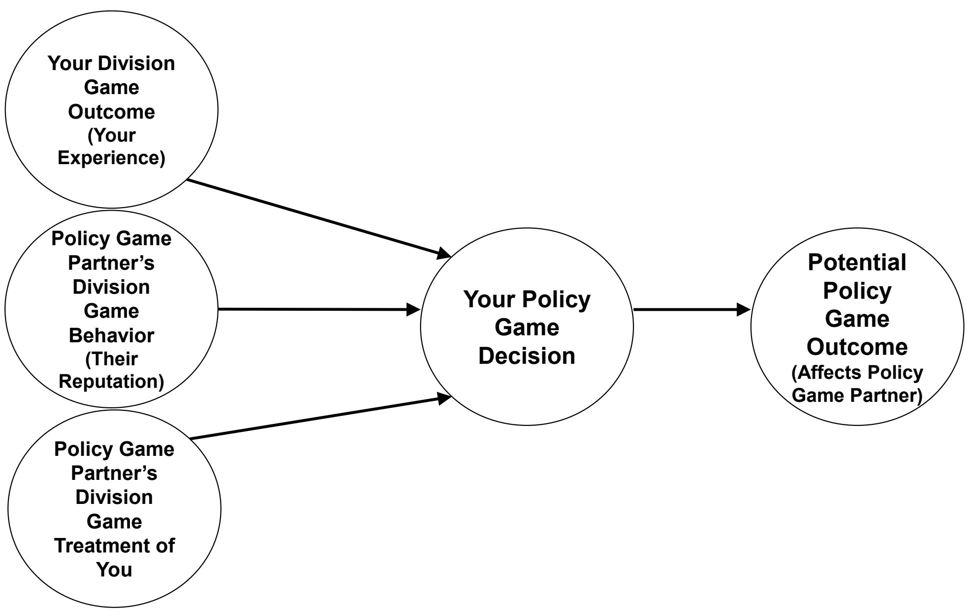

Results

In the Division Game, 63% of the Receiver participants (whose decisions were not implemented) proposed the equal 20/20 division, 14% proposed the inequitable 30/10 division, and 23% proposed the 40/0 division. In the Policy Game, these participants were assigned to the role of Player A and were equally and randomly assigned to each of the three Division Game outcomes. Additionally, they were randomly assigned at equal rates to one of the three policy conditions shown in Table 2 which varied in the payoffs to them while holding constant the payoffs to their Policy Game partner (Player B).

We start our analysis by examining the effect of receiving either unfair division (10/30 or 0/40) in the Division Game on choices in the Policy Game. Pooling across policy and pairing conditions, participants who received either unfair division in the Division Game were 11 points less likely to choose the cooperative policy (72% versus 83%) that would improve the payoffs of their partner by 20 tokens (rather than simply retaining their initial allocation) compared to participants who were given a fair split of 20 of the 40 allocated tokens.

We confirm this result with formal statistical analysis in Table 3. Participants who got 0 tokens in the Division Game were 13 points ($p < .01$) less likely to endorse the cooperative policy compared to participants who received a fair split of 20 tokens (Column 1). Similarly, participants who received 10 tokens were 10 points less likely to do so. These differences are indistinguishable from one another. When controlling for policy proposal payoffs (Column 2), the effect of receiving either unfair division is to reduce the likelihood of choosing the cooperative policy by 9 points ($p < .01$). Notably, as expected, participants were much more likely to adopt the cooperative policy when failing to do so was personally costly: They were 39 and 45 points more likely to propose adopting the policy when it netted them 10 or 20 tokens, respectively.

Finally, we turn to our key manipulation: How did the pairing condition affect Policy Game decisions, specifically the relationship between experience in the Division Game and support for the cooperative policy? We do so by estimating a regression model in which we add to the column 2 specification from Table 3 by also allowing the effect of the Division Game outcome to vary by the partner condition. This analysis appears in Column 3. Per these results, there is no statistically significant relationship between the Division Game outcome and support for the policy when the Division Game was played with a different person, even in the Duplicate Person condition. The coefficients for Division Game Got 0 and Got 10 are between -3 and -4 points, or less than half of the -9 points from the column (2) specification.

Instead, it appears that the effect of either unfair allocation in the Division Game is much larger when both games were played with the same person. The coefficient for Division Game Got 0 x Same Person is -13 ($p<.05$) and Division Game Got 10 x Same Person is -11 points ($p=.07$), implying that participants in the Same Person condition were 15 points (got 0 tokens, -2.5-12.6, $p < .01$, got 10 tokens, -.4.2-10.6) less likely to vote for a policy when they received 0 or 10 tokens in the Division Game relative to when they were allocated 20 tokens by their partner. Proportionally the effect of receiving 0 tokens is 6 times larger in the Same Person condition than in the Duplicate Person condition and 3.5 times larger when the Division Game outcome is 10 tokens. These results support the Partner History effect that receiving an unfair division undermines subsequent cooperation. Thus, "losing" in the division of resources only undermined subsequent cooperation when both tasks were done with the same individual.

Last, we examine two potentially interesting sources of heterogeneity in these patterns. First, we estimate the effect of playing both games with the same person on voting in specific policy conditions. In accordance with our theoretical expectations, participants reacted more to their partner's selfish behavior when it was less costly to do so, that is when they did not receive any net tokens from choosing the policy in the Policy Game. We arrived at this conclusion after re-estimating the column 3 specification from Table 3 after partitioning the data into two groups:

cases where the Policy Game nets 0 and cases where it nets 10 or 20 tokens for the participant (Complete analysis appears in Appendix Table A2). In the Net 0 Policy condition, participants who received 0 tokens from their partner in the Same Person condition were 35 points less likely to support the policy ($p < .01$) than participants who received an equal 20/20 split and the coefficient on the interaction of Division Game Got 0 x Same Person is -29 points ($p<.05$). By contrast, participants in the Same Person condition were no less likely to support the policy when they received 0 tokens from their partner in the Net 10 or Net 20 condition compared to when they received an equitable division. In this subsample, the coefficient on the interaction of Division Game Got 0 x Same Person is -3 points ($p=.57$). We find similarly larger effects for getting only 10 tokens in the division game in the Net 0 compared to Net 10 and Net 20 conditions. The aggregate pattern we find is therefore that spillover effects are largest when individuals have the least to lose by refusing to cooperate in the policy game. Spillover in the Same Person condition appears to be most pronounced when there is no net benefit to an individual's personal benefit to cooperating, even though overall welfare is improved by doing so.

Figure 2, in which we present the full breakdown of support for the policy proposal by partner condition (top panel is Same Person, bottom panel is Duplicate Person), Decision Game outcome (Got 0, 10, or 20), and Policy (Net 0, 10, or 20) confirms this regression analysis. It is only for the Net 0 (leftmost part of both panels) that the effect of Decision Game outcomes affects support for the policy, and the effect is much stronger in the Same Person condition (top panel).

Additionally, we also explored whether there were differences in responses to unfair Division Game outcomes in general or in the Same Person condition by whether a participant's own proposed division (not implemented) was fair or not. To do so, we looked at two subsamples, one where the participant proposed an even fair split (620 participants) and the second where they proposed a highly unfair split of 40/0 (230 participants). Complete analysis appears in Appendix Table A3. We estimate that among those who proposed an even split the effect of receiving 0 tokens (rather than 20) in the Division Game on support for the policy is -3 points (n.s.) in the Duplicate Person condition and -17 points in the Same Partner condition ($p<.01$). The equivalent estimates are -5 point (n.s.) and -17 points ($p<.05$) among those who proposed the exact same unfair split in the Division Game. We therefore find little evidence that individuals respond differently to unfair Division Game behavior by whether they themselves proposed an unfair split in the Division Game.

Figure 2: Experiment 1. Proportion proposing to adopt policy, by pairing condition, Division Game outcome, and Policy Game net policy payoff.

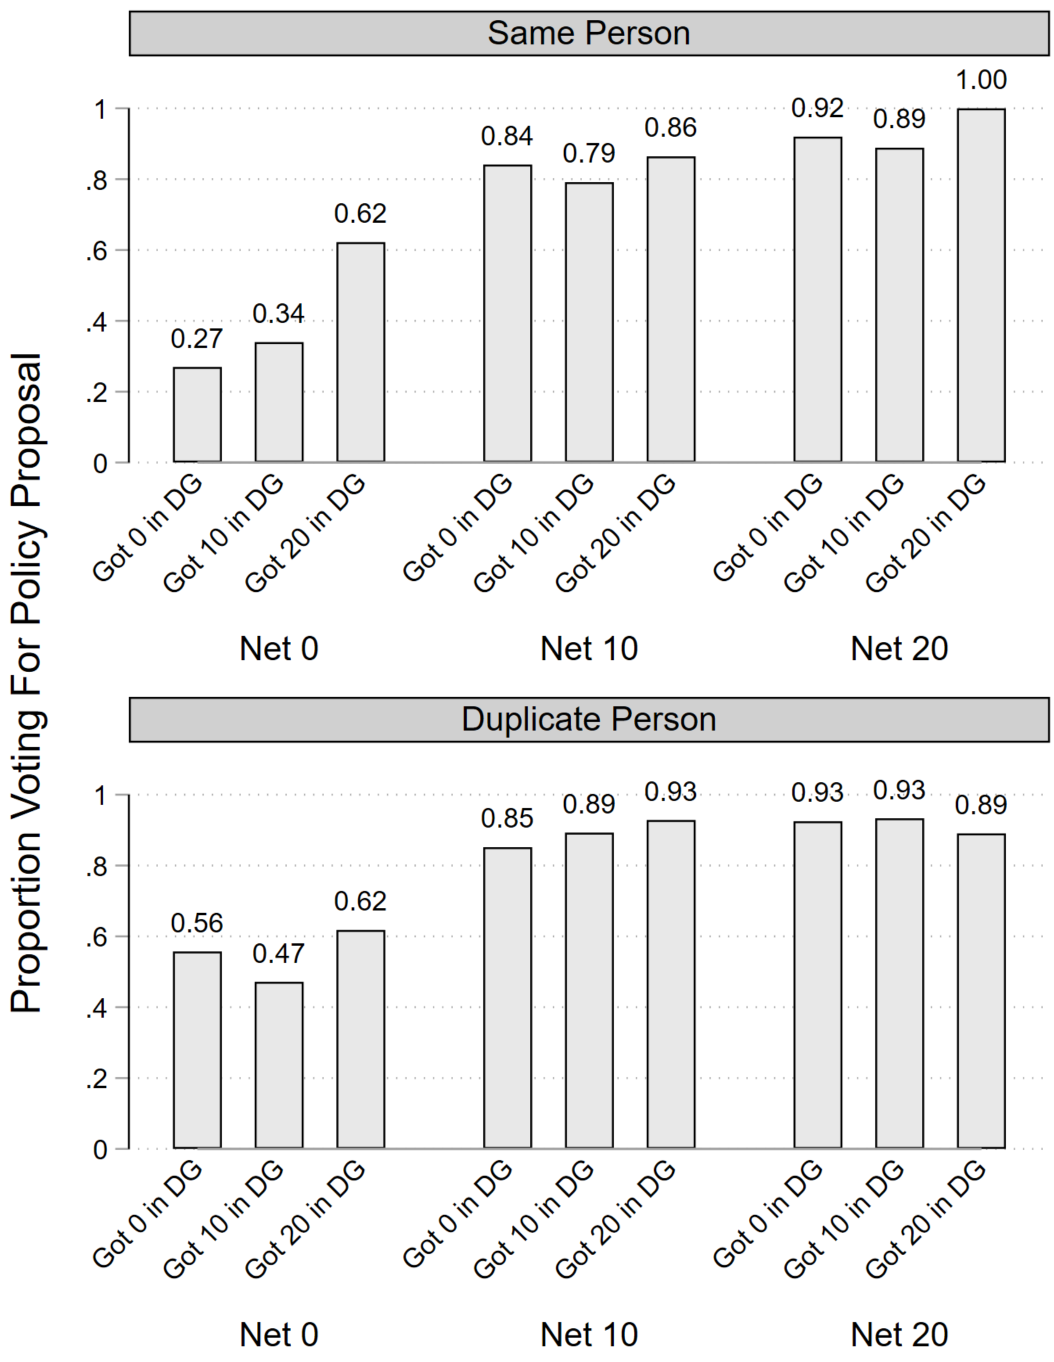

Horizontal Axis is Division Game Outcome (Top) and Policy Game Net Payoff to Player (Bottom)

Table 3: Experiment 1. Participants are less likely to support the policy when they receive an unfair division in the Same Person condition.

	(1)	(2)	(3)
	DV: Propose to adopt policy in Policy Game (1 = yes)		
Division Game Got 0	-0.13***	-0.09***	-0.03
	[0.03]	[0.03]	[0.04]
Division Game Got 10	-0.10***	-0.09***	-0.04
	[0.03]	[0.03]	[0.04]
Policy Game Nets 10		0.39***	0.39***
		[0.03]	[0.03]
Policy Game Nets 20		0.45***	0.45***
		[0.03]	[0.03]
Partner is Same Person			0.01
			[0.04]
Division Game Got 0 x Same Person			-0.13**
			[0.06]
Division Game Got 10 x Same Person			-0.11*
			[0.06]
Constant	0.83***	0.54***	0.53***
	[0.02]	[0.03]	[0.04]
R^2	0.02	0.22	0.23
N	987	987	987

*** *p*<.01, ** *p*<.05, * *p*<.10. The table reports unstandardized OLS coefficients with robust standard errors in brackets. Got 20 tokens in the Division Game, Policy Game Net 0, and Duplicate Person are the excluded categories.

Discussion

Taken together, these results show that receiving an unfair split in the Division Game reduced subsequent support for a cooperative policy in the Policy Game, but only when players interacted with the same individual in both games even though the different individual's Division Game behavior was identical to their partner's. This Partner History Effect appears to be driven by deviations from equity, as any unfair division reduced support for a policy in the Same Partner condition. However, there is limited evidence of spillover when a division is made by a different individual than the individual who benefits from the policy being adopted in the Policy Game. We also find evidence that being the loser in the division of resources undermines cooperation much more when individuals only break even on policy, meaning that choosing not to cooperate is costless. Thus, cooperation is not always undermined by receiving an unfair division, it is undermined by receiving an inequitable allocation from a person who stands to gain from that subsequent cooperation.

Experiment 2: Group History vs. Personal History

In Experiment 1, participants either interacted with the same individual for both games or a different individual in each game, but this set aside groups which are an important feature of social life. While receiving an inequitable allocation from a "stranger" did not affect support for a policy, participants did not receive any information about the individual other than their past behavior. Importantly, their Policy Game partner was in no way affiliated with the individual responsible for the Division Game allocation. But many allocation decisions are made by members of groups and these decisions are often influenced by group membership.

To examine the role of groups, we modified our experimental framework to include group membership. Before playing the Division Game, participants completed a clicking task in which they had 5 seconds to click their mouse as many times as possible. Participants were told that this task would be used to determine their group affiliation and that they would be put into groups with other participants who had a similar "clicking pattern" to them.[4] This was done to instantiate an attachment to the group, consistent with previous work using a "minimal groups paradigm" (Tajfel & Turner, 1979). Participants received 1 token for each 4 clicks that they completed, rounded down to the nearest cent. This pattern of clicks was used to place participants into one of four groups for the subsequent stages of the experiment. We used a procedure that assigned participants to groups such that there was no correlation between the

[4]Participants were assigned to groups with the following equation: (number of clicks mod 4). By doing so, there is a correlation of 0 between the number of times participants clicked and their group assignment.

number of times that they clicked and the group to which they were assigned. All participants were told that they were in Group A given that the specific labels for groups were arbitrary.

Next, participants played the Division Game, as described in Experiment 1, with two notable differences. First, participants were always told that they were playing with a member of Group B, i.e. an out-group member. Second, given the similarity in the effects from Experiment 1 of receiving 0 or 10 tokens, Deciders could either take all 40 tokens (leave none for their partner) or take 20 tokens (leave 20 for their partner, an equitable division). We adopted this approach for reasons of statistical power.

After the Division Game, participants played the Policy Game. Before voting on the policy, participants learned four pieces of information. They learned:

1. Whether their partner in this game was also their partner in the Division Game (a specific member of Group B)
2. Whether their partner was a member of Group B or Group C
3. Whether their partner was the Decider in the Division Game
4. Their partner's earnings in the Division Game

We randomly assigned participants to play the Policy Game with the same member of Group B who they played the Division Game with (same manipulation in Experiment 1), a different member of Group B who played the Division Game with a member of Group D (not the player's group), or a member of Group C (an altogether new group) who played the Division Game with a member of Group D. For the last two pairings, we chose Group D partnership to avoid invoking any in-group/out-group dynamics beyond what was directly experienced by participants in the Division Game. As in Experiment 1, regardless of a participant's pairing they were always told about an individual whose behavior in the Division Game exactly matched that of their own partner in the Division Game to allow us to isolate the causal effect of pairing condition by holding the General Experience Effect constant. That is, regardless of group membership, a new partner always duplicated the behavior of their Division Game partner.

After learning this information, players indicated their support for the proposed policy and it was randomly determined whose decision was implemented.

We again recruited participants from Amazon Mechanical Turk. These participants were paid 50 cents for completion and a bonus payment based on their behavior in each stage. We again completed our recruitment in two stages. In the first wave, we recruited participants to take the role of the Decider in the Decision Game. As before, these participants made decisions for all three policies in the Policy Game and were paid accordingly.

In the second wave, we recruited participants (n = 1,776) who were not Receivers in the Division Game and took the role of Player A in the Policy Game. These participants were randomly assigned to a Division Game allocation (0 tokens or 20 tokens which correspond to real decisions made by a human player recruited earlier), a policy (Net 0, Net 10, or Net 20), and a pairing condition (*Same Person*: play both stages with the same person, *Same Group*: play both stages with different persons from Group B, or *Duplicate Person*: play the Division Game with a member of Group B and the Policy Game with a member of Group C). We under-sampled the Net 20 condition because participants in Experiment 1 voted for that policy at near ceiling levels. The distribution of subjects in policy conditions was approximately 40% in Net 0, 40% in Net 10, and 20% in Net 20 (see Table A4 for sample size by condition). Again, participants answered comprehension questions to confirm that they understood the game. Participants answered a total of ten comprehension questions and 81% of them answered at least 9 correctly.

Figure 3: Causal models for effects in Experiment 2
Panel A: General Experience Effect (Duplicate Person condition)

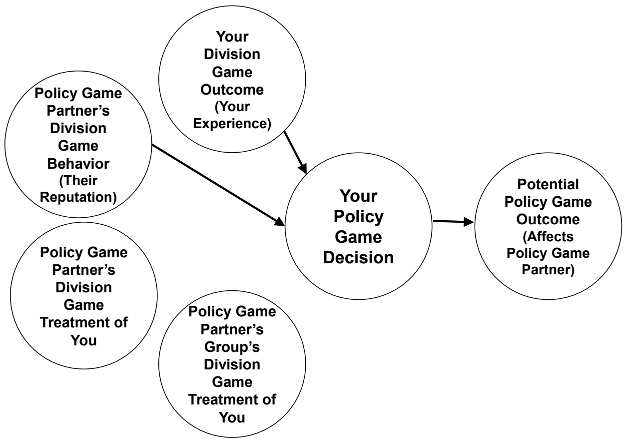

Panel B: Partner History Effect (Same Person condition)

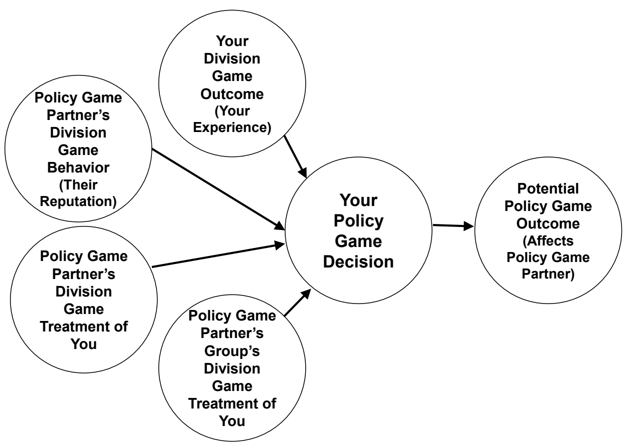

Panel C: Group History Effect (Same Group condition)

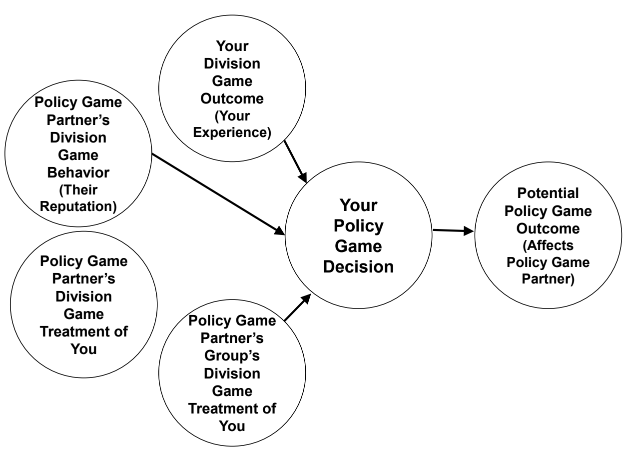

Figure 3 shows the causal relationships of interest in Experiment 2. The introduction of groups creates a situation in which an individual who is the Receiver in the Division Game could have a personal history with either an individual or a group. In the absence of interacting with the same individual or different members of the same group repeatedly, an individual can draw on their experience in the Division Game and how their partner behaved in the Policy Game to determine whether to support the policy (Figure 3 Panel A). This is the General Experience Effect. As was the case in Experiment 1, interacting with a specific individual repeatedly creates a shared history that may cause spillover between games, but it also creates a history with an outgroup (Figure 3 Panel B). Thus, it could be the case that receiving an inequitable allocation from one member of a group creates spillover if an individual interacts with different members of the same group in the Division Game and the Policy Game (Figure 3 Panel C). This Group History Effect is distinct from the Partner History Effect as it does not involve interaction with the *same* outgroup member. As before, we can therefore compare outcomes across conditions to isolate the Group and Partner history effects.

Results

As with Experiment 1, we focus our analysis on participants who were Receivers in the Division Game because our key treatment variable is the outcome in the Division Game as generated by another player. In the Division Game, 65% of these Receiver participants (whose decisions were not implemented) proposed the equal 20/20 division and 35% proposed the inequitable 40/0 division. We focus here on estimating the interactive effect of pairing condition and receiving an unfair division. Following our earlier specification, we interact the pairing condition with an indicator for receiving an unfair outcome to predict support for the policy

(Table 4). We use the Duplicate Person condition as a baseline for comparison and control for the effect of each policy in the Policy Game.

As expected, the effect of Division Game outcome is largest, 18 points (-6-12, $p < .01$) in the Same Person condition and smallest, 6 points ($p = .06$) in the Duplicate Person condition. The effect of an inequitable outcome in the Same Group condition falls in between at 10 points (-6-4, $p < .01$). Put differently, the effect of the Same Group condition was 56% as large as the effect of the Same Person condition, while the effect of the Duplicate Person condition was 33% as large. The effect of the Same Person condition is statistically distinguishable from the effect in the Duplicate Person condition ($p < .01$), though it was only qualitatively larger than the Same Group condition ($p < .09$). The difference between being matched with a Duplicate Person compared to a person from the Same Group was not statistically significant ($p = .30$).

These results show that interacting with the same individual is the largest driver of spillover effects. We do, however, find weaker spillover effects also occur when individuals interact with two different members of the same group, although this effect is not statistically distinct from when an individual interacts with a member of one out-group in the Division Game and a member of a different out-group in the Policy Game. We acknowledge greater statistical imprecision in comparison involving the Same Group condition.

As before, we also examine spillover effects by policy condition to understand the effects of variation in the cost of failing to cooperate. These results appear in columns (2) and (3) of Table 4 where we subset our analysis first to cases where the policy nets the player 0 tokens and then to cases where it nets the player 10 or 20 tokens. We observe spillover in all of the pairing conditions for the Net 0 policy condition. Specifically, participants in the Same Person condition were 38 points ($p < .01$) less likely to vote for the policy when they received an unfair allocation compared to an equitable allocation. Participants in the Same Group and Duplicate Person condition were 17 points ($p < .01$) and 15 points ($p < .05$) to vote in favor of the policy, respectively. The effect for the Same Person condition is distinguishable from the others at p<.05.

Figure 4: Experiment 2. Proportion proposing to adopt policy, by pairing condition, Division Game outcome, and Policy Game net policy payoff.

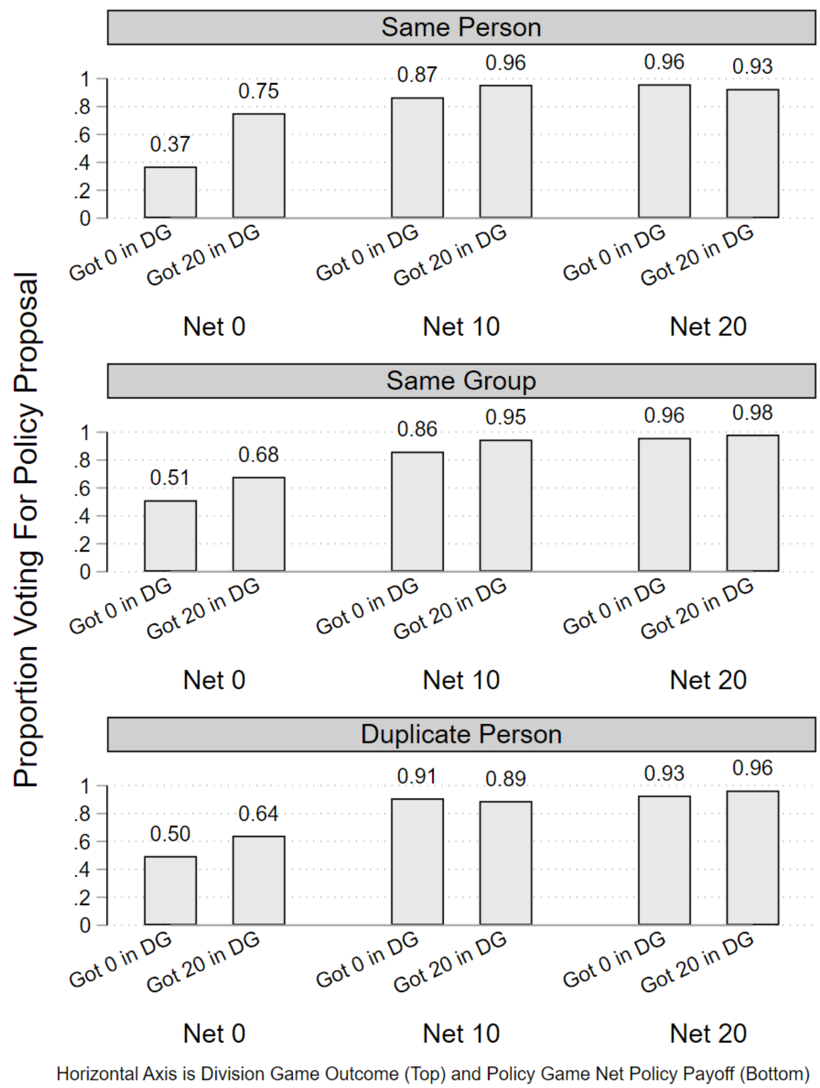

Table 4: Experiment 2. Effect of Division Game outcome and pairing condition on voting in the Policy Game

	(1)	(2)	(3)
	All Cases	Net Policy Payoff 0	Net Policy Payoff 10 or 20
	DV: Propose to adopt policy in Policy Game (1 = yes)		
Division Game Got 0	-0.06*	-0.15**	-0.00
	[0.03]	[0.06]	[0.03]
Policy Game Nets 10	0.32***		
	[0.02]		
Policy Game Nets 20	0.38***		0.05***
	[0.02]		[0.02]
Partner is Same Person	0.06**	0.11*	0.03
	[0.03]	[0.06]	[0.03]
Partner is Same Group	0.04	0.04	0.04*
	[0.03]	[0.06]	[0.03]
Division Game Got 0 x Same Person	-0.12***	-0.24***	-0.05
	[0.04]	[0.09]	[0.04]
Division Game Got 0 x Same Group	-0.04	-0.02	-0.06
	[0.04]	[0.09]	[0.04]
Constant	0.60***	0.64***	0.90***
	[0.03]	[0.04]	[0.02]
R^2	0.19	0.07	0.01
N	1,776	691	1,085

*** p<.01, ** p<.05, * p<.10. The table reports unstandardized OLS coefficients with robust standard errors in brackets. Got 20 tokens in the Division Game, Policy Game Net 0 (Column 1) and Net 10 (Column 3), and Duplicate Person are the excluded categories.

For the Net 10 and 20 conditions, however, participants in the Same Group and Same Person condition were estimated to be 5 to 6 points less likely to vote for the policy when they received an inequitable allocation, but these estimates are not statistically significant. The point estimate in the Duplicate Person condition is 0. As with the first experiment and in accordance with our expectations, these results demonstrate that spillover effects are most pronounced in situations that involve group conflict when the personal benefit from cooperation is comparatively low.

The raw display of this voting behavior in Figure 4, where we further partition the Net 10 and Net 20 cases, further confirms this interpretation. Clear evidence of spillover effects is only present in the Net 0 policy cases and is greatest when interacting with the same person, following by members of the same group.

In the appendix (See Table A5) we also test whether there are differences in these effects by a participant's own Decision Game proposal. Unlike with Experiment 1, in Experiment 2 we find that spillover effects are larger when participants proposed fair splits in the Division Game than when their own proposals were unfair. It is most notable in the Same Person condition, where we estimate that the effect of an unfair Division Game outcome is to reduce support for the policy by 22 points when the participant proposed a fair split but only 9 points when the participants proposed a selfish split. We do not have a theoretical explanation for the difference we find in the results between Experiment 1 and 2.

Discussion and Conclusion

In a pair of controlled experiments, we examined whether being the loser in an allocation of resources subsequently reduced cooperation. In both experiments, we find evidence in favor of a Partner History effect: Interacting with the same individual causes reduced support for a policy after receiving an inequitable allocation of resources. In Experiment 2, we also find evidence of a Group History effect. The effect of interacting with two different out-group members was approximately half as large as the effect of interacting with the same out-group member. However, we do find considerable spillover effects in the Net 0 policy condition when participants interacted with any out-group member (the same individual, two different members of the same out-group, or individuals who were members of different out-groups). Taken tougher, these findings demonstrate that interacting with the same individual is the largest driver of spillover effects, although out-group animosity also appears to result in spillover leading to less cooperation when some individuals are only weakly better off from their provision.

These results support the account that institutions that sequester conflictual private goods from the provision of cooperative goods avoid harmful spillover. This aligns with past work arguing in favor of localizing responsibility for the provision of local public goods or the management of common pool resources (Ostrom, 1990). More broadly, this work also suggests that believing that certain issues, for example racial equality (Norton & Sommers, 2011), are zero-sum games may not only be damaging in those policy domains, but could also spill over into domains where there is a potential for cooperation (Pareto-improvement).

The experimental design that we employed allows us to draw clear causal inferences about the effect of pairing conditions on spillover. However, this abstraction comes with a trade off in that it sets aside potentially relevant factors that are present in real world political interactions. For example, we ignore the possibility that separating related tasks that have different payoff structures may silo relevant expertise, which may be important for the possibility of identifying cooperative outcomes in the first place. Additionally, participants were not informed about the existence of the Policy Game until they completed the Division Game. Knowing about the second stage could, of course, reduce the prevalence of inequitable divisions by forward-looking actors seeking to avoid the costly spillovers that we observed. Further, in Experiment 2, we employed a minimal group paradigm that established a weak connection between participants and their group. This design choice allowed us to rule out potential confounds that using actual groups would have introduced into our experimental design. Further, by not repeating the game multiple times, we eliminated the possibility of turn taking, which has the potential to mute spillover effects.

Nonetheless, this work has identified an important connection between dividing limited resources and cooperative tasks from which everyone can benefit: That losing in the division of resources causes harmful spillover into Pareto-improving opportunities that benefit the winner. This finding lays the foundation for future investigation into whether real world institutions, prevent spillover from occurring and if they do not, what kinds of changes to institutional rules could eliminate spillovers that hurt the welfare of all individuals.

References

Ahn, T.-K., Ostrom, E., Schmidt, D., Shupp, R., & Walker, J. (2001). Cooperation in PD gamesy. *Public Choice,* 106(1), 137-155.

Albert, M., Güth, W., Kirchler, E., & Maciejovsky, B. (2007). Are we nice (r) to nice (r) people?—an experimental analysis. *Experimental Economics,* 10(1), 53-69.

Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *science,* 211(4489), 1390-1396.

Balliet, D., Wu, J., & De Dreu, C. K. (2014). Ingroup favoritism in cooperation: a meta-analysis. *Psychological bulletin,* 140(6), 1556.

Bednar, J., Chen, Y., Liu, T. X., & Page, S. (2012). Behavioral spillovers and cognitive load in multiple games: An experimental study. *Games and Economic Behavior,* 74(1), 12-31.

Bicskei, M., Lankau, M., & Bizer, K. (2016). Negative reciprocity and its relation to anger-like emotions in identity-homogeneous and-heterogeneous groups. *Journal of Economic Psychology,* 54, 17-34.

Bokemper, S., Descioli, P., & Kline, R. (2019). Unfair Rules for Unequal Pay. *Journal of Experimental Political Science*, 6(3), 180-191.

Burnham, T. C. (2015). Public goods with high-powered punishment. *Journal of Bioeconomics*, 17(2), 173-187.

Camerer, C. F., & Thaler, R. H. (1995). Anomalies: Ultimatums, dictators and manners. *Journal of Economic perspectives,* 9(2), 209-219.

Capraro, V., & Marcelletti, A. (2014). Do good actions inspire good actions in others? *Scientific reports,* 4(1), 1-6.

Cason, T. N., Savikhin, A. C., & Sheremeta, R. M. (2012). Behavioral spillovers in coordination games. *European Economic Review,* 56(2), 233-245.

Chang, H. I., Dawes, C. T., & Johnson, T. (2018). Political inequality, centralized sanctioning institutions, and the maintenance of public goods. *Bulletin of Economic Research,* 70(3), 251-268.

d'Adda, G., Capraro, V., & Tavoni, M. (2017). Push, don't nudge: Behavioral spillovers and policy instruments. *Economics Letters,* 154, 92-95.

Engelmann, D., & Fischbacher, U. (2009). Indirect reciprocity and strategic reputation building in an experimental helping game. *Games and Economic Behavior,* 67(2), 399-407.

Fowler, J. H., & Kam, C. D. (2007). Beyond the self. *The Journal of politics,* 69(3), 813-827.

Jordan, J. J., Hoffman, M., Nowak, M. A., & Rand, D. G. (2016). Uncalculating cooperation is used to signal trustworthiness. *Proceedings of the National Academy of Sciences,* 113(31), 8658-8663.

Knez, M., & Camerer, C. (2000). Increasing cooperation in prisoner's dilemmas by establishing a precedent of efficiency in coordination games. *Organizational behavior and human decision processes,* 82(2), 194-216.

Landa, D. & Duell, D. (2015), Social Identity and Electoral Accountability. *American Journal of Political Science*, 59(3), 671-689.

Milgrom, P. R., North, D. C., & Weingast, B. R. (1990). The role of institutions in the revival of trade. *Economics & Politics,* 2(1), 1-23.

Misch, A., Paulus, M., & Dunham, Y. (2021). Anticipation of future cooperation eliminates minimal ingroup bias in children and adults. *Journal of Experimental Psychology: General*. 150(10), 2036–2056.

Norton, M. I., & Sommers, S. R. (2011). Whites see racism as a zero-sum game that they are now losing. *Perspectives on Psychological science,* 6(3), 215-218.

Nowak, M. A., & Roch, S. (2007). Upstream reciprocity and the evolution of gratitude. *Proceedings of the royal society B: Biological Sciences,* 274(1610), 605-610.

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*: Cambridge university press.

Sheffer, Lior, and Peter John Loewen. (2019). Accountability, framing effects, and risk-seeking by elected representatives. *Political Research Quarterly* 72(1), 49-62.

Tajfel, H., & Turner, J. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worchel (Eds.), *The social psychology of intergroup relations* (pp. 33-48). Monterey, CA: Brooks/Cole.